# Applied Neuroscience

## Columbia Science Honors Program

**Spring 2017**

## Learning and Memory

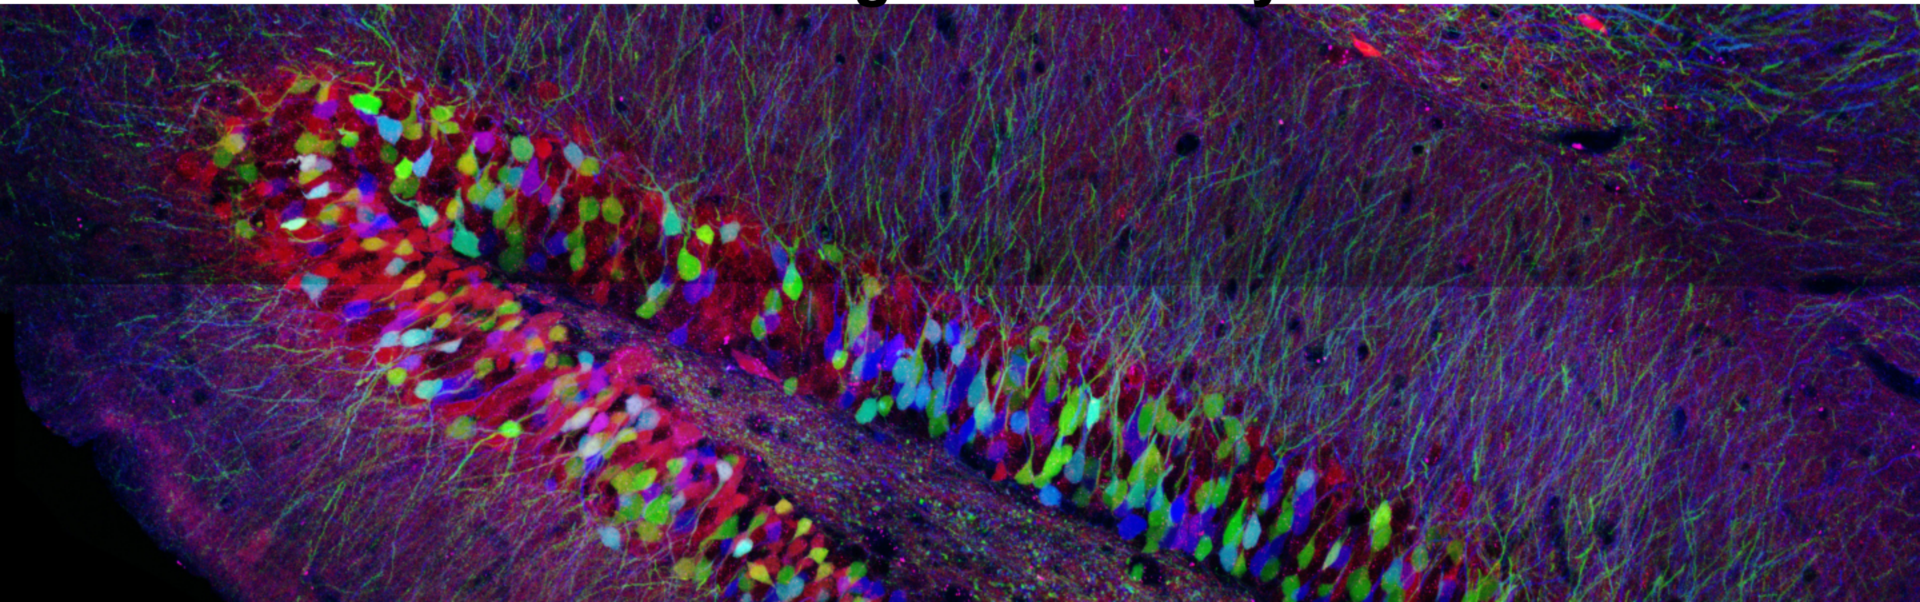

# Learning and Memory

**Objective:** Computational Models of Memory

**Agenda:**

1. Learning and the Brain
   - Introduction to Neural Networks
2. Memory
   - Structure and Function of Memory
   - Memory Disorders
3. How to Read Scientific Papers
   - *Creating a False Memory in the Hippocampus*

     Steve Ramirez et al. Science 2013; 341: 387-391

# Learning and the Brain

Dogs and cats are both furry animals with four legs and many other shared traits. Why, then, is it easy to distinguish between them?

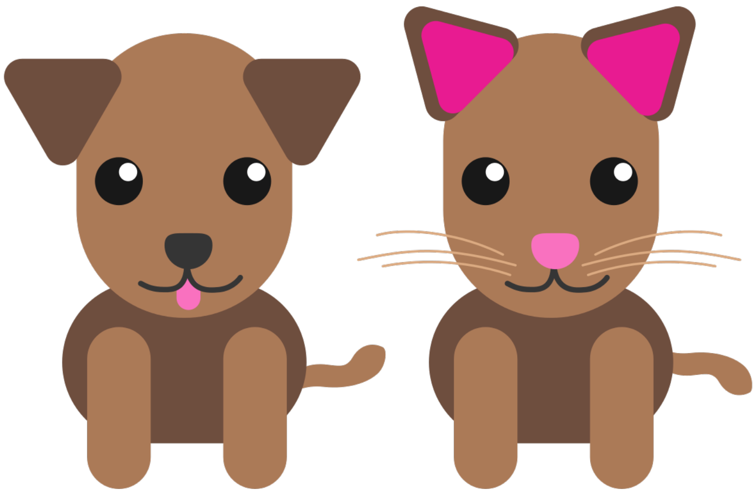

As young people, we’re told which animals we observe are dogs and which are cats. Fairly quickly, we stop needing new examples. Our learning is powerful enough to classify a new animal as a dog or a cat, even when it doesn’t look particularly similar to one we’ve seen before. It turns out that computers learn similarly.

# Learning and the Brain

A **supervised learning** algorithm attempts to model a function to relate inputs to outputs. It uses known examples to learn this relationship.

When building a supervised learning model to distinguish whether an image is of a dog or a cat, what should the inputs for the examples be?

A. The fur and eye colors of the dogs and cats
B. The lengths and weights of dogs and cats
C. Numerical data representing images of dogs and cats

# Learning and the Brain

A **supervised learning** algorithm attempts to model a function to relate inputs to outputs. It uses known examples to learn this relationship.

When building a supervised learning model to distinguish whether an image is of a dog or a cat, what should the inputs for the examples be?

A. The fur and eye colors of the dogs and cats
B. The lengths and weights of dogs and cats
C. **Numerical data representing images of dogs and cats**

# Learning and the Brain

Suppose you have access to 100,000 images of dogs and cats that you can use to build a supervised learning model that distinguishes between dogs and cats.

After using images as examples to **train** (or teach) the model, you'll want to use images to **test** the model; that is, to determine if the model is actually successful at identifying if the image is a of a dog or cat.

# Learning and the Brain

What would be a reasonable way to select your images for training and testing?

A. Use the dog pictures for training and the cat pictures for testing
B. Use the cat pictures for training and the dog pictures for testing
C. Split the images randomly into two sets: one for training and one for testing
D. Use all of the images for training and testing

# Learning and the Brain

What would be a reasonable way to select your images for training and testing?

A. Use the dog pictures for training and the cat pictures for testing
B. Use the cat pictures for training and the dog pictures for testing
**C. Split the images randomly into two sets: one for training and one for testing**
D. Use all of the images for training and testing

# Learning and the Brain

A good supervised learning model predicts the outputs of unobserved inputs using knowledge of the outputs of observed inputs. The ability to make successful predictions on unobserved inputs from observed data is called **generalization.**

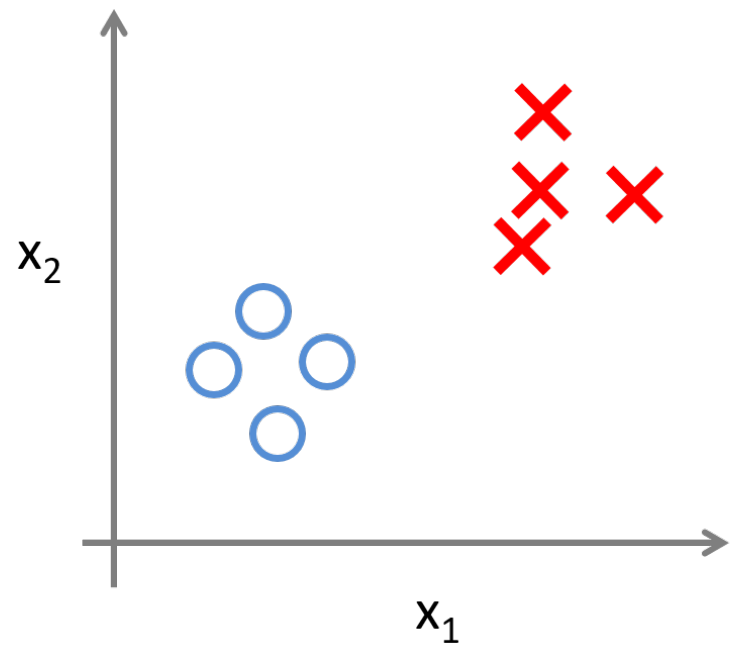


For any observed data there are an infinite number of functions that pass through all input-output pairs. The “best” function is *not* necessarily the one which fits all the observed data, but instead the one that generalizes well.

# Learning and the Brain

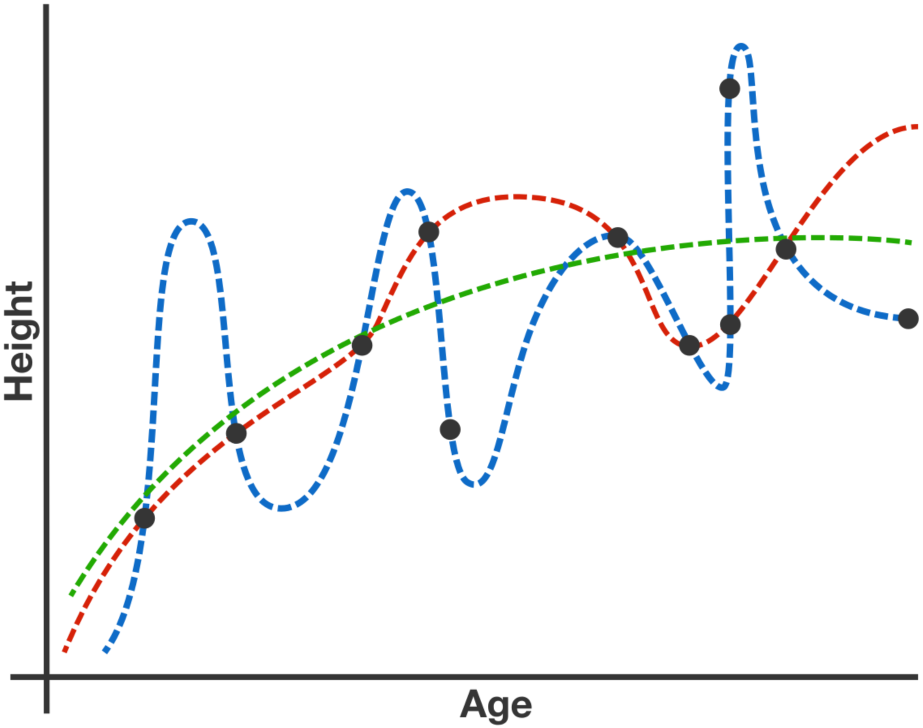


A. Blue Model
B. Green Model
C. Red Model

You are training a height-prediction model using observed inputs of age and outputs of height, shown as points in the graph above. Which of these three functions drawn is likely to be the best model?

# Learning and the Brain

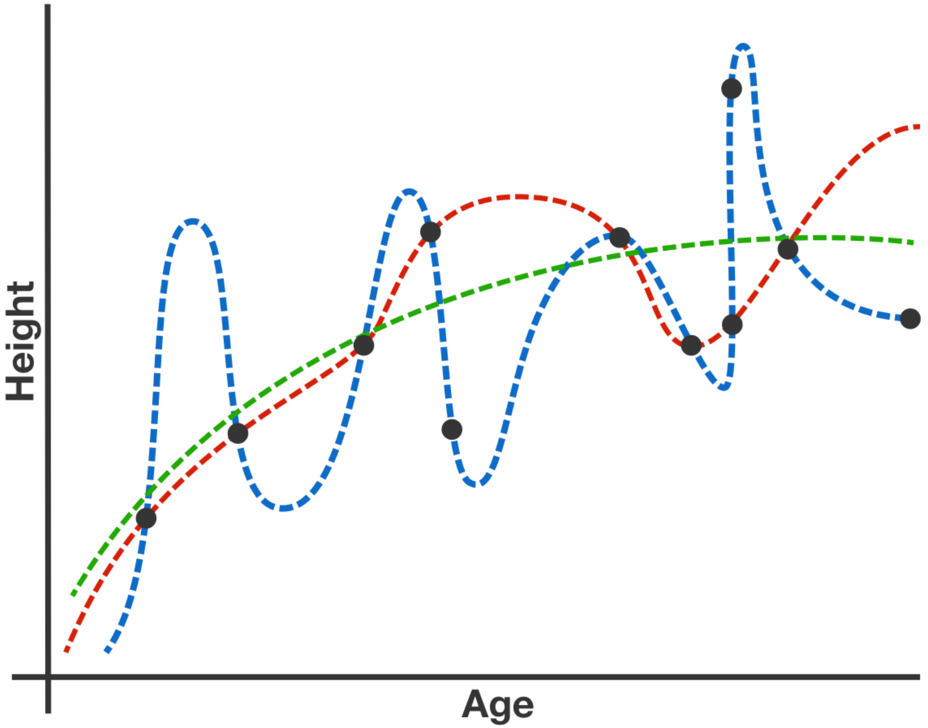


A. Blue Model
**B. Green Model**
C. Red Model

You are training a height-prediction model using observed inputs of age and outputs of height, shown as points in the graph above. Which of these three functions drawn is likely to be the best model?

# Learning and the Brain

In distinguishing between dogs and cats, we wanted to classify an image into discrete categories with no numerical relationship; i.e., we cannot say a *dog* is *2 times cat.* This type of problem is called a **classification** problem.

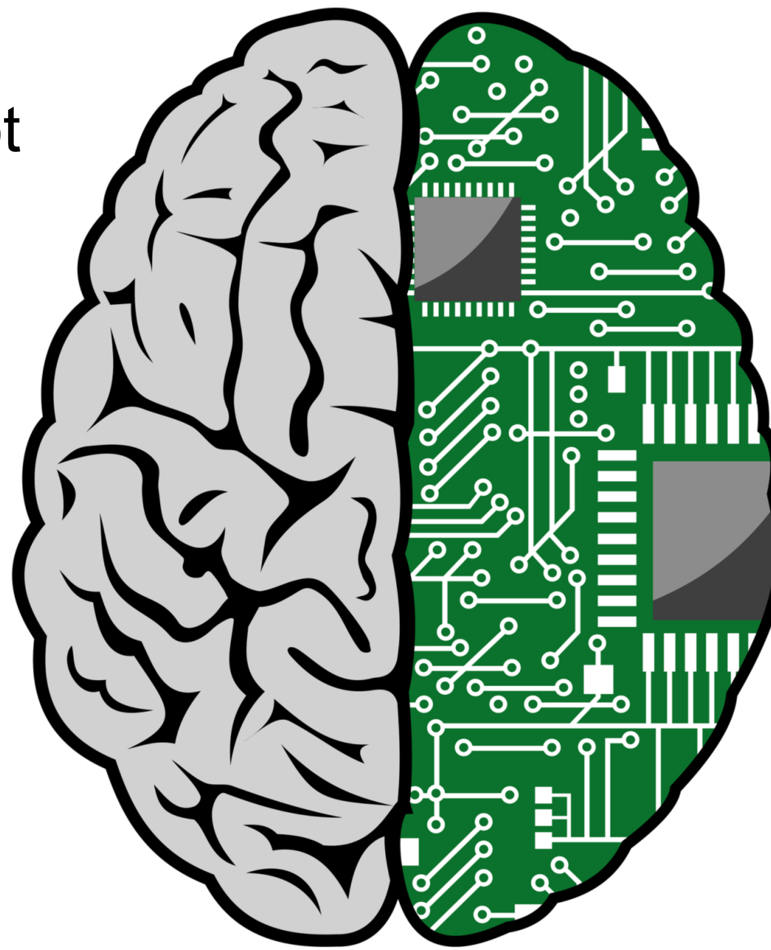

In the last question, we found a function to relate an input to a numerical output (height). These outputs have a clear numerical relationship. This type of problem is known as a **regression** problem.

**Artificial neural networks** can be used to solve both types of problems.

# Introduction to Neural Networks

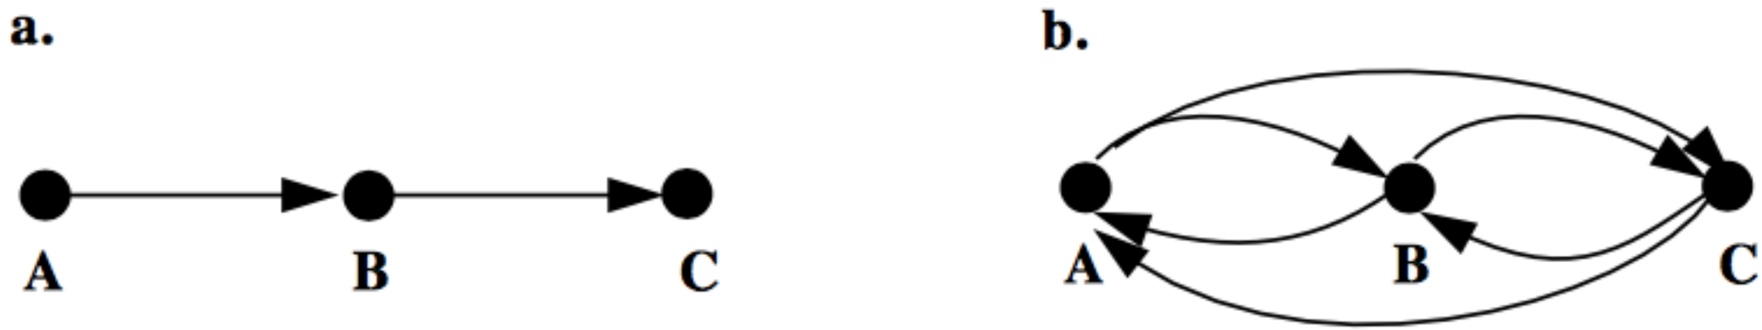


**A. Feed-forward network:** Information flows directly from one layer of neurons to the next without feedback
*In the brain, the existence of such network connectivity is rare (an example of it is between the retina and the LGN)*

**B. Recurrent network:** Information is connected with feedback
*In the brain, these networks are ubiquitous.*

# Role of Feedback in Neural Networks

***What difference does feedback make in a neural network?***
Feedback is a way to create a memory from memory-less components. More specifically, we create a persistent state through non-persistent parts with **feedback.**

***"Reverberating Activity Loops"*** *were proposed by Rafael Lorente de No and Donald Hebb. The idea is that excitation among units in a circuit **enables** excitation to persist **beyond the duration of exciting stimulus.***

***Short-term memory*** *uses changes in activity*
***Long-term memory*** *uses changes in synaptic connections (**Long-term potentiation (LTP)** is used for memory acquisition)*

# Learning and the Brain

Of the following three learning problems, how many should be treated as **regression** problems?

- Identifying which zip-code digits have been written on an envelope
- Predicting the total number of points scored by two teams in a basketball game
- Determining the *probability* that a given person will get heart disease


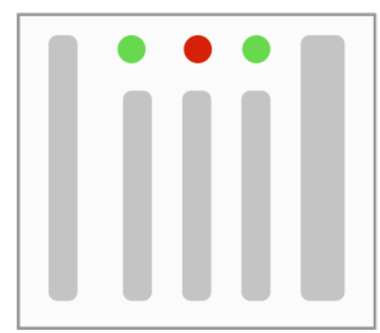

# Learning and the Brain

Of the following three learning problems, how many should be treated as **regression** problems?

- Identifying which zip-code digits have been written on an envelope
- **Predicting the total number of points scored by two teams in a basketball game**
- **Determining the *probability* that a given person will get heart disease**


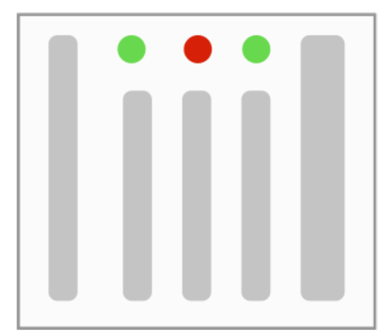

# Learning and the Brain

We have focused on supervised learning problems. However, there is another type of learning: **unsupervised learning.**

Unsupervised learning attempts to determine relationships between inputs without using an example outputs (such as “dog” or “cat”). Which of the following would **not** be a good fit for unsupervised learning?

- Determining possible friendship matches from interest profiles
- Predicting someone’s gender from their name
- Grouping news articles about similar topics

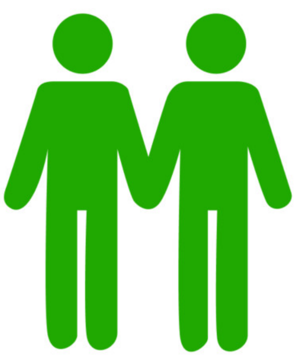

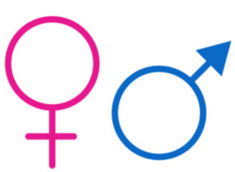

# Learning and the Brain

We have focused on supervised learning problems. However, there is another type of learning: **unsupervised learning.**

Unsupervised learning attempts to determine relationships between inputs without using an example outputs (such as “dog” or “cat”). Which of the following would **not** be a good fit for unsupervised learning?

- Determining possible friendship matches from interest profiles
- **Predicting someone’s gender from their name**
- Grouping news articles about similar topics

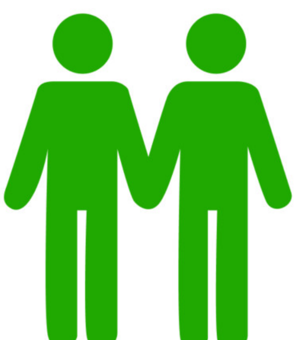

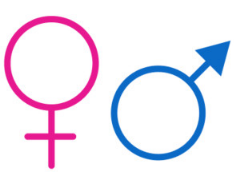

# Learning and the Brain

SINGULARITY
SUMMIT

# Learning and the Brain

- Batch Learning
- Online Learning

When training a learning model, there are two main processes that can be used with respect to how the training data is handled: batch learning and online learning.

In **batch learning**, the model learns from batches of data- often the entire training set at once. In **online learning**, the model learns from data processed sequentially over time, as it becomes available. Which type of learning does the human brain use?

# Learning and the Brain

- Batch Learning
- **Online Learning**

When training a learning model, there are two main processes that can be used with respect to how the training data is handled: batch learning and online learning.

In **batch learning**, the model learns from batches of data- often the entire training set at once. In **online learning**, the model learns from data processed sequentially over time, as it becomes available. Which type of learning does the human brain use?

# Learning and the Brain

Based on your intuition, which of the following is **not** an advantage of online learning over batch learning?

- It is more efficient for data storage
- It allows for gradual improvement over time
- It makes performance evaluation of the learning model easier

# Learning and the Brain

Based on your intuition, which of the following is **not** an advantage of online learning over batch learning?

- It is more efficient for data storage
- It allows for gradual improvement over time
- **It makes performance evaluation of the learning model easier**

# Learning and the Brain

Artificial neural networks can use both online learning and batch learning.

**Neurons** are the simplest units of the computation in the human brain. While the actual dynamics of a neuron’s computation are complex, a simplified view of them is that they **integrate and fire.** That is, a neuron performs a computation with its inputs, and then fires if that computation passes a certain threshold.

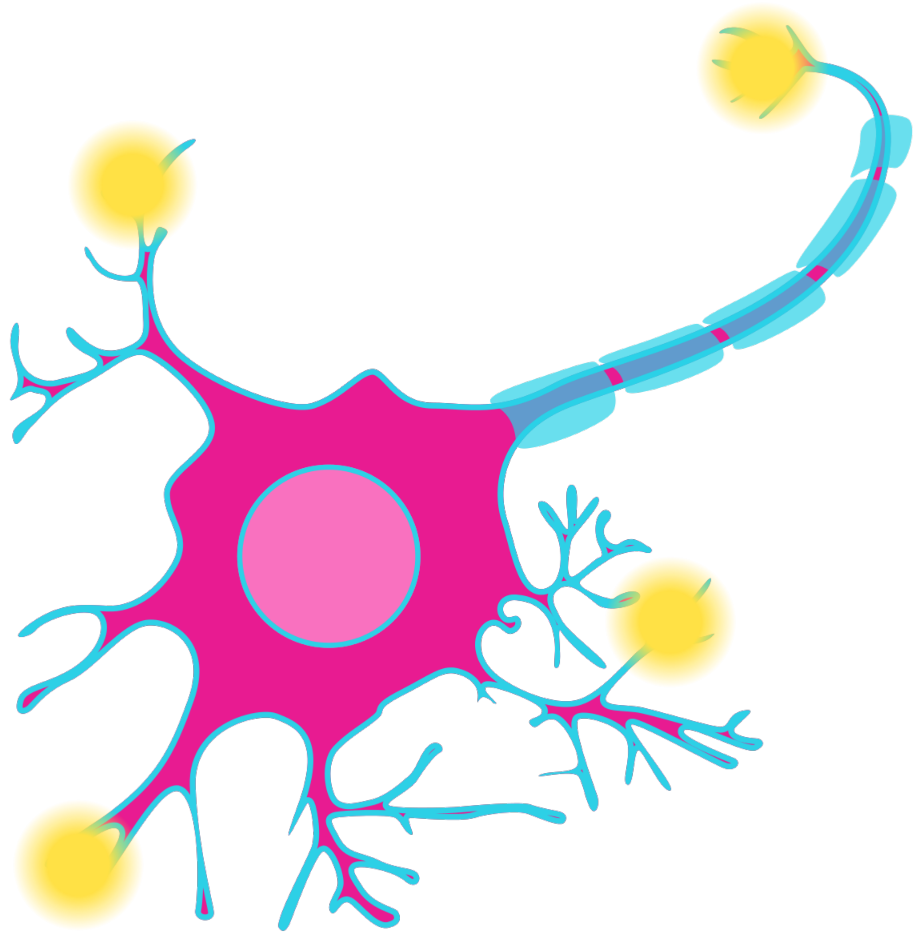

# Learning and the Brain

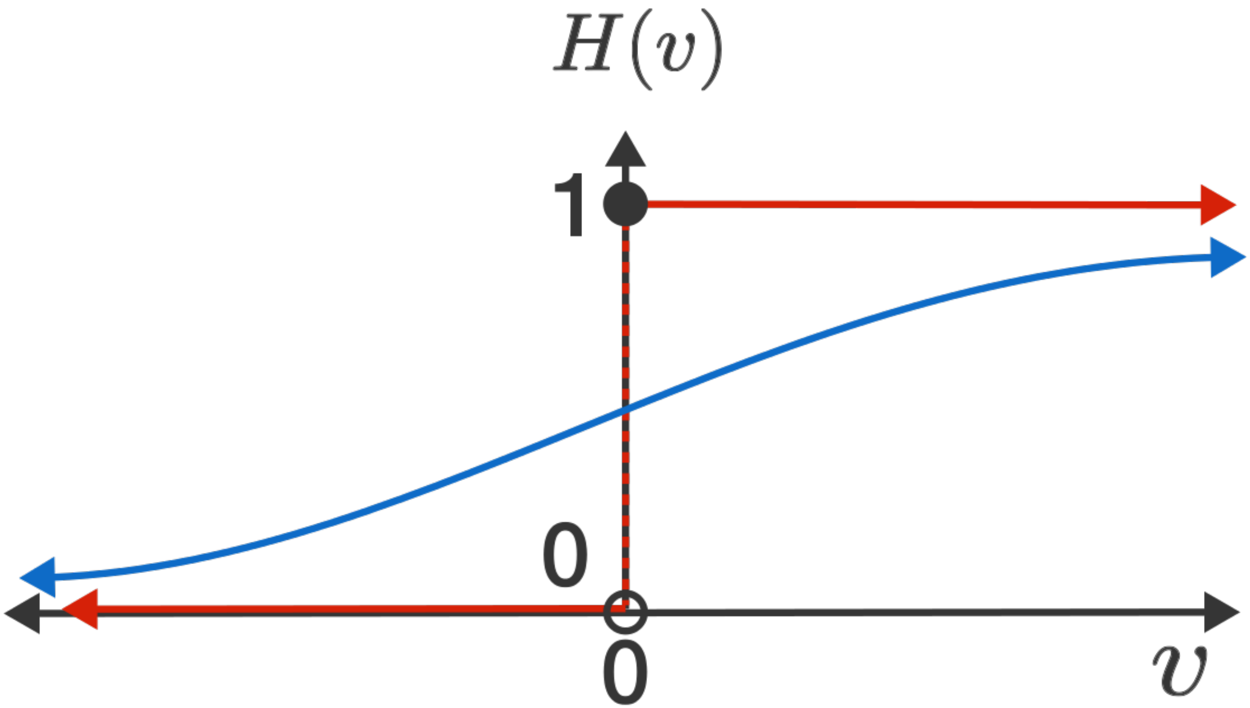


- Red Model
- Blue Model

In an artificial neural network, the inputs are combined as a weighted sum into a single value, *v.* Then, an activation function, *H(v)*, is applied to determine whether or not the neuron fires.

What activation function would best model a physical neuron?

# Learning and the Brain

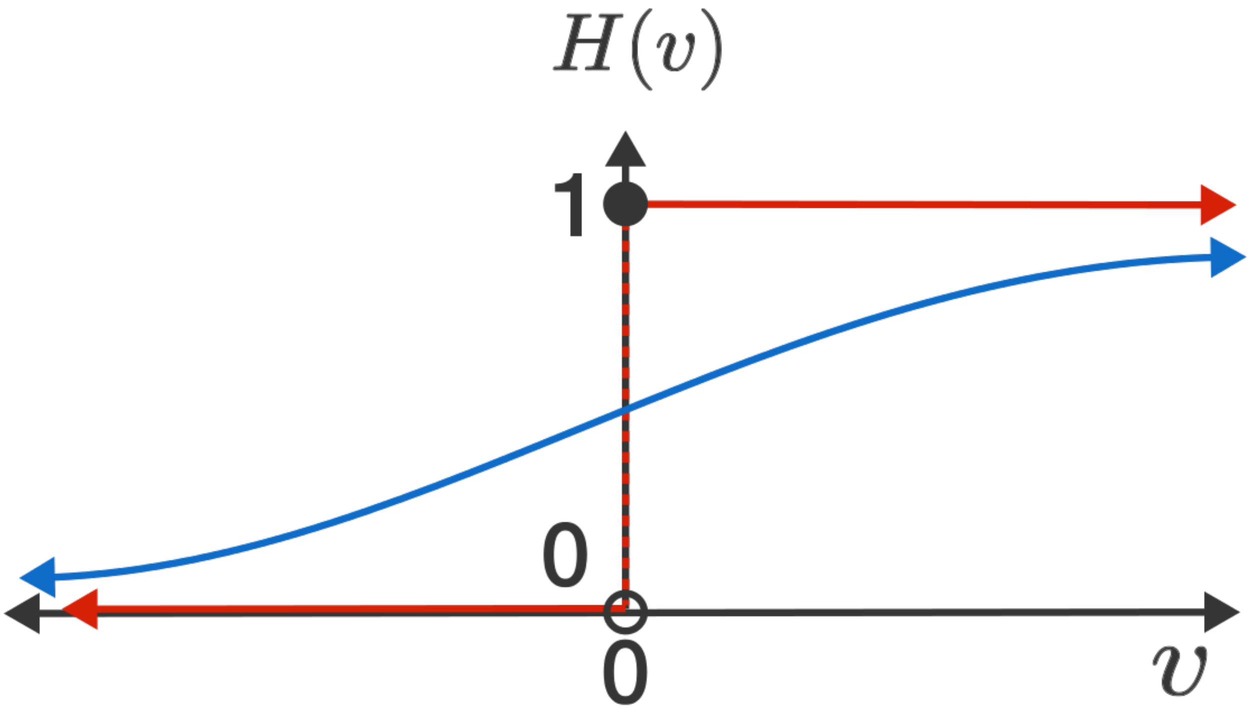


- **Red Model**
- Blue Model

In an artificial neural network, the inputs are combined as a weighted sum into a single value, *v.* Then, an activation function, *H(v)*, is applied to determine whether or not the neuron fires.

What activation function would best model a physical neuron?

# Learning and the Brain

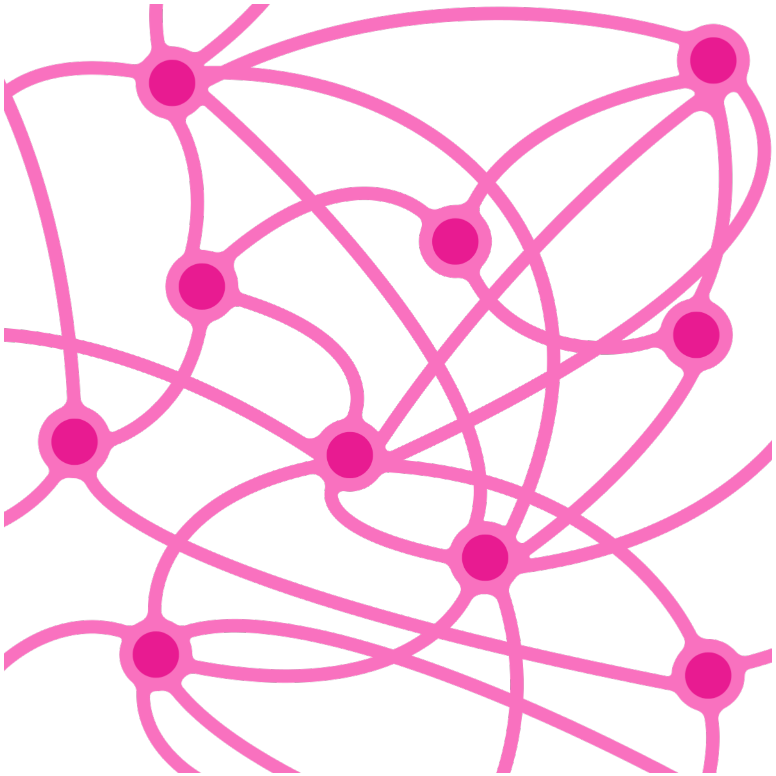

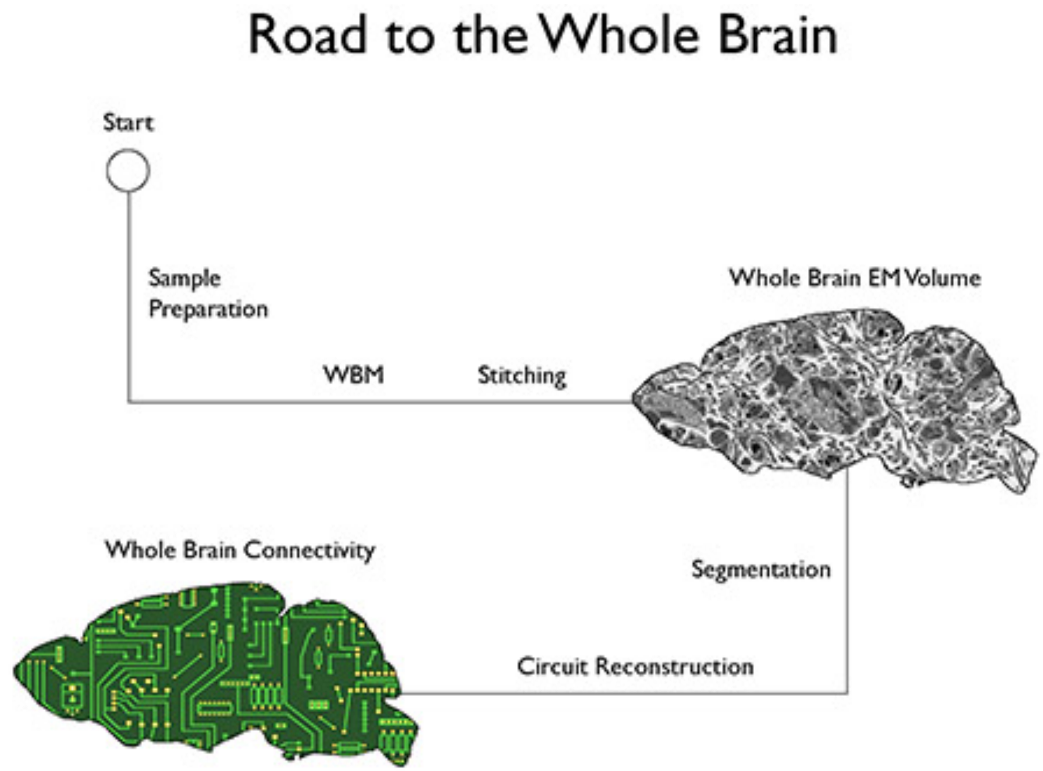


Connections (or synapses) between neurons are used to pass information from the outputs of some neurons to the inputs of other neurons. Not every neuron is connected to every other neuron, and certain neurons have stronger connections than others.

# Learning and the Brain

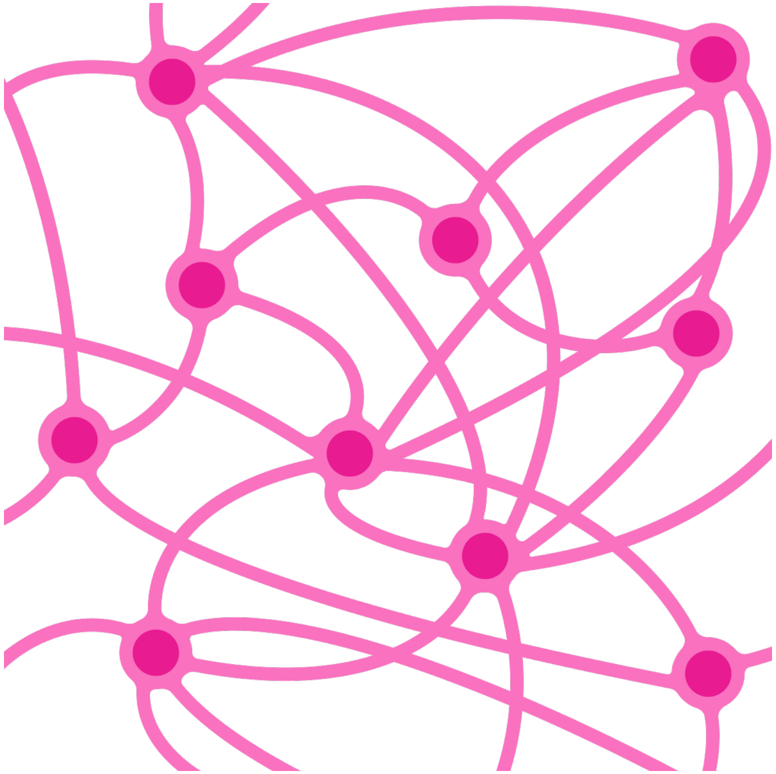

If there are approximately $\mathbf{10^{11}}$ neurons in the human brain and between $\mathbf{10^{14}}$ and $\mathbf{10^{15}}$ synapses, which of these values is a reasonable estimate for the average number of connections per neuron?

- 10
- 1,000
- 100,000
- 1,000,000

# Learning and the Brain

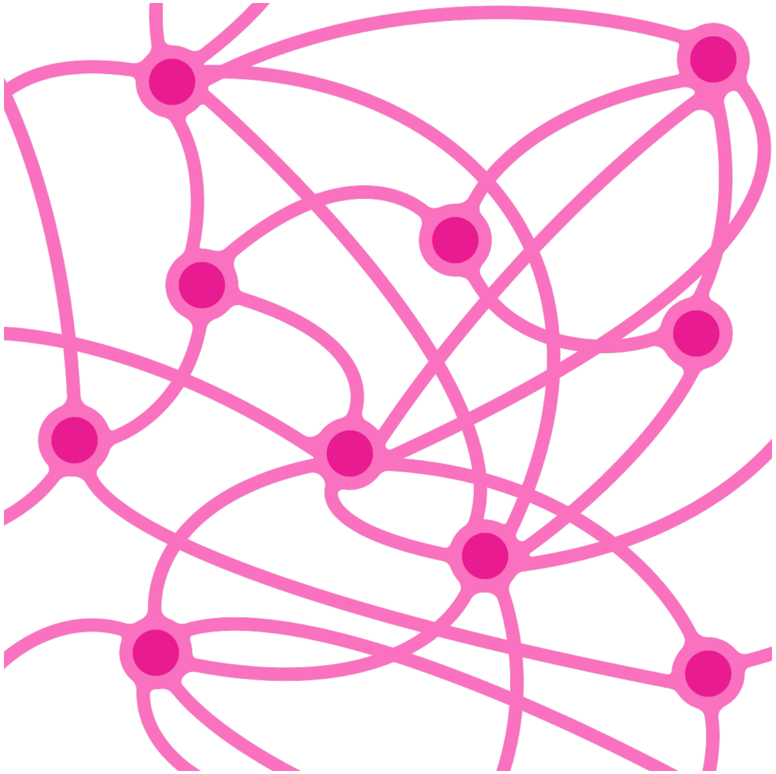

If there are approximately **$10^{11}$** neurons in the human brain and between **$10^{14}$** and **$10^{15}$** synapses, which of these values is a reasonable estimate for the average number of connections per neuron?

- 10
- **1,000**
- 100,000
- 1,000,000

# Learning and the Brain

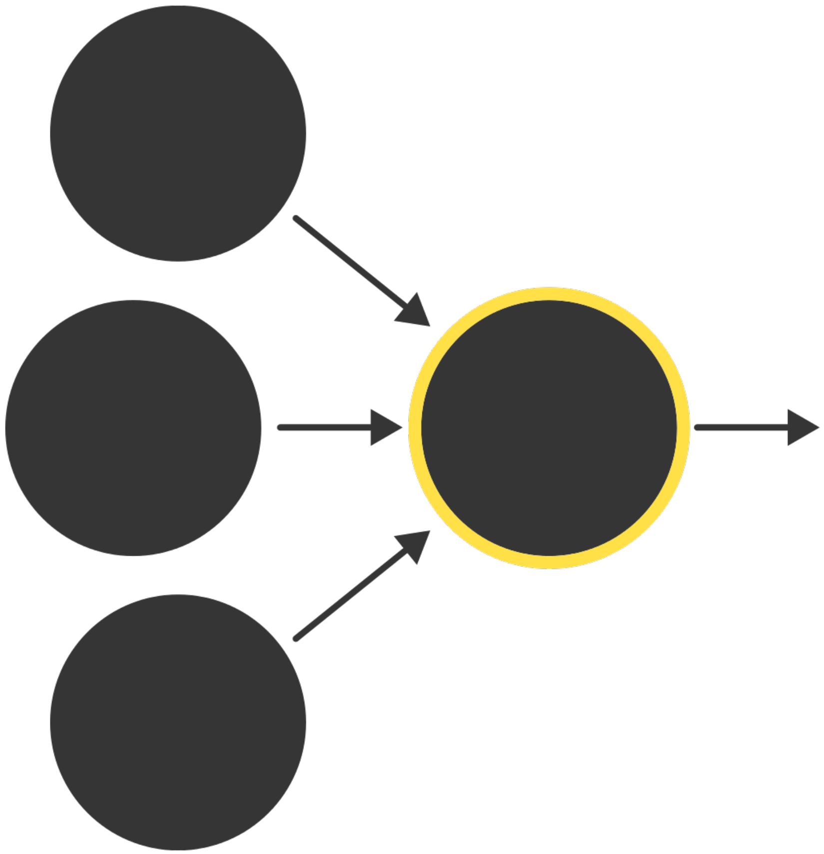

By adjusting which connections exist and how strong they are, the human brain is able to learn a huge variety of complex functions. Thus, a computational model of the human brain should include simple computational units (like neurons) which are connected to one another (like synapses). If the model can learn to adjust the strengths of those connections appropriately, it may be able to approach the power of the human brain.

# Introduction to Memory

## *What is memory?*

**Memory:** changes in the activity or connectivity of neural systems that are triggered by stimuli or brain states and ***persist over a duration longer than triggering events***

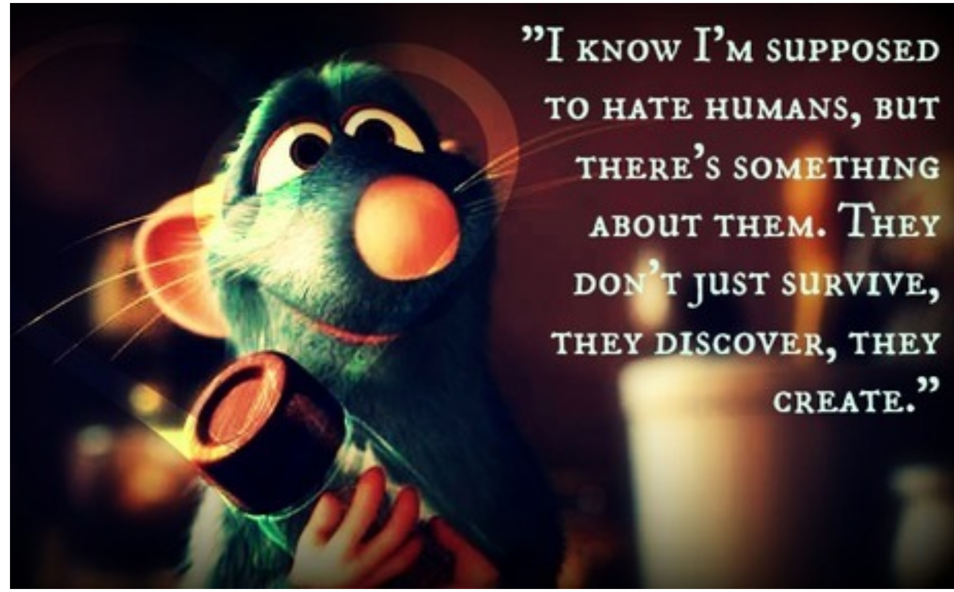


## *Why do we need memory?*

Memory serves an **adaptive role:**

1. We learn from our experiences
2. We generalize faster
3. We can make predictions

***Learning, Interference, and Prediction***

Memory research uses:

1. Neuroscience
2. Computer Science
3. **Information Theory:** studies the quantification, storage, and communication of information

# Introduction to Memory

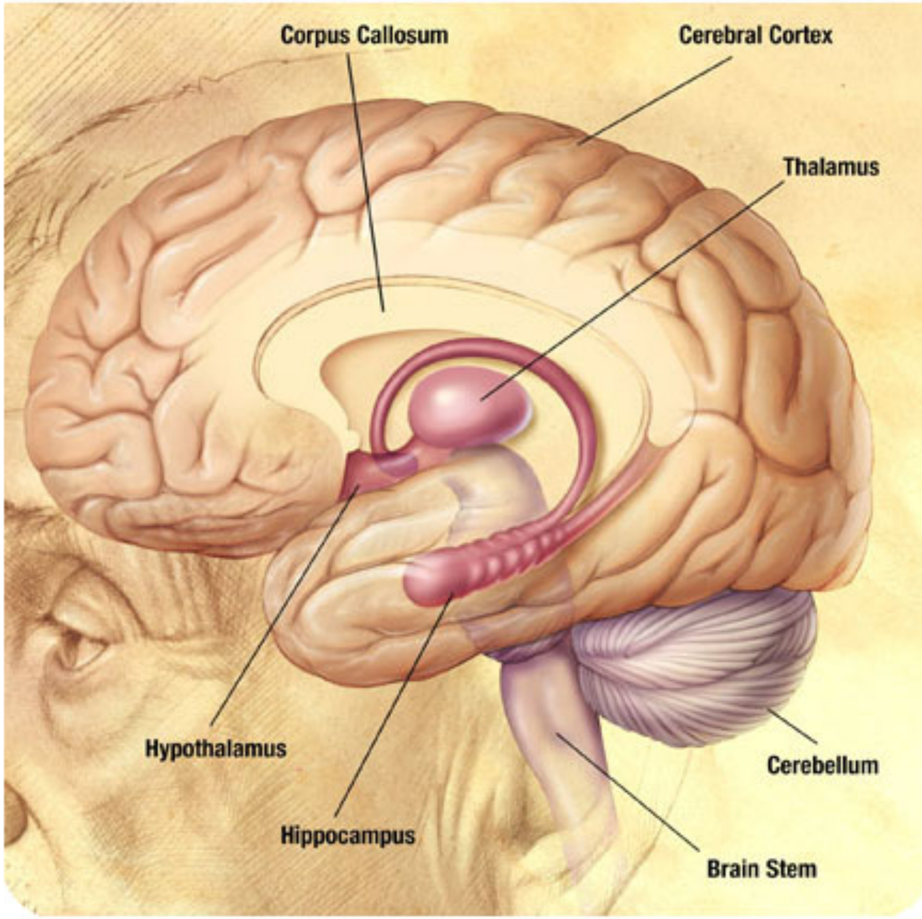


**Limbic System:** controls emotions and instinctive behavior (includes hippocampus and parts of cortex)

**Thalamus:** receives sensory and limbic input and sends to cerebral cortex

**Hypothalamus:** maintains homeostasis and controls internal clock of body

**Hippocampus:** where short-term memories are converted to long-term memories

# Introduction to Memory

Where is memory?

**Multiple Brain Regions** are involved in encoding memory as shown below by fMRI studies

**fMRI:** uses MRI technology that measures brain activity by detecting changes in associated with blood flow. Cerebral blood flow and neuronal activation are coupled.

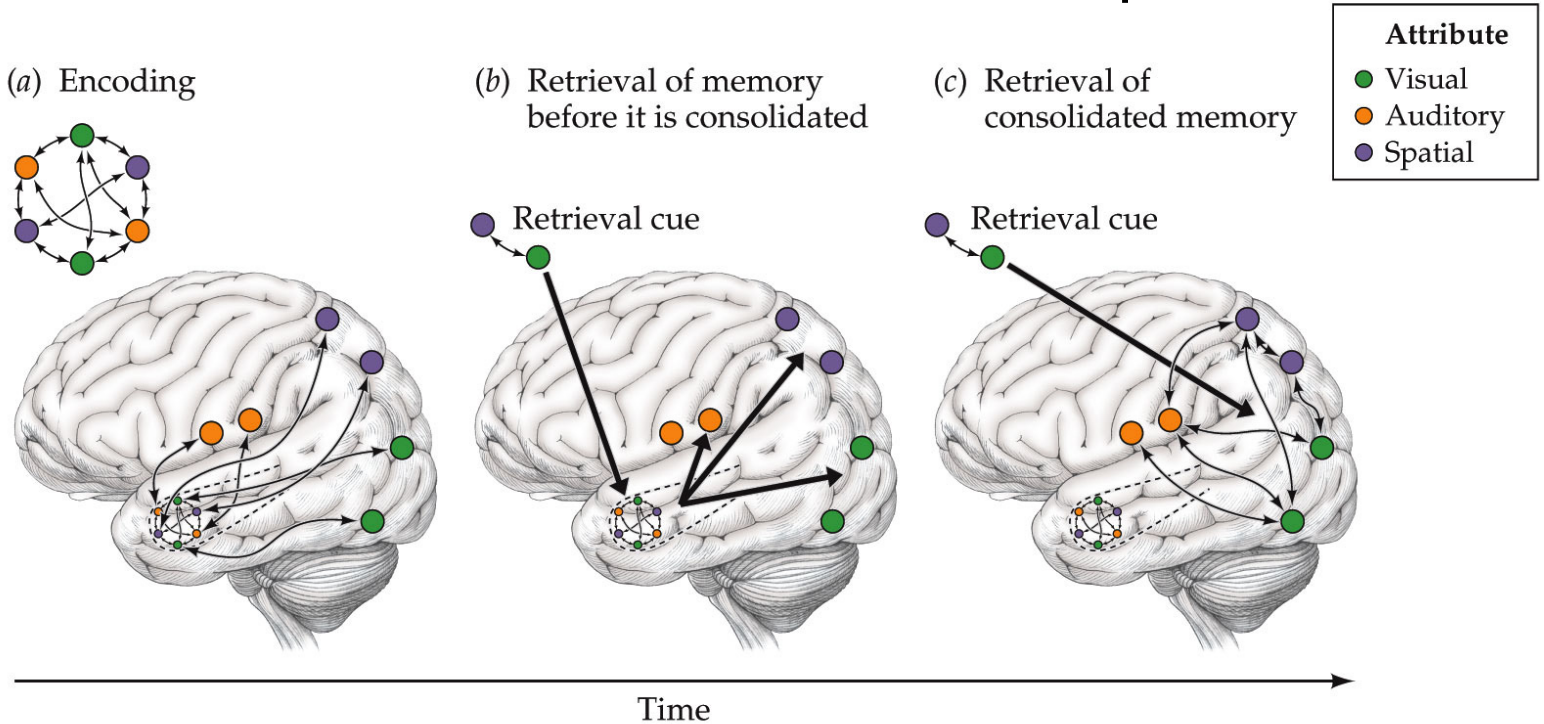

# Types of Memory

- Models of memory include a sequence of three stages: **sensory, short-term and long-term memory**
- Different types of memory have their own mechanism of action
- This sequential model is called modal, multi-store or the Atkinson-Shiffrin model (developed in 1968)

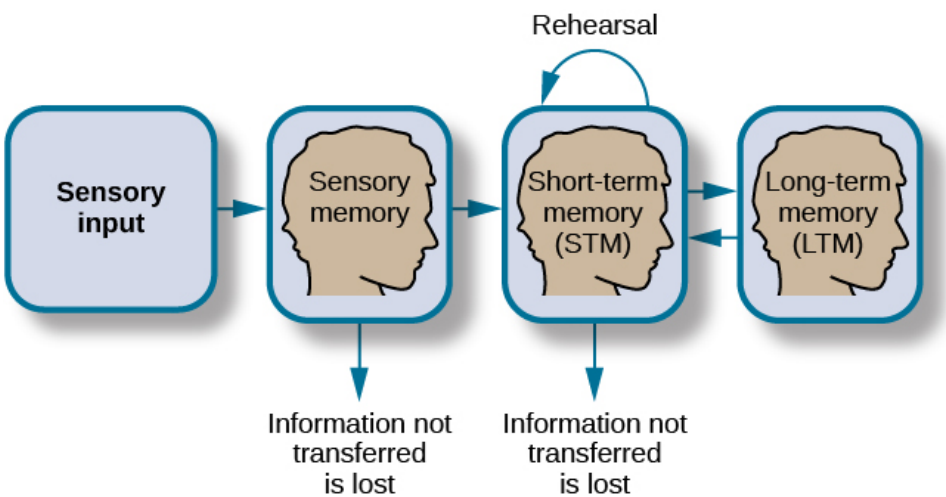


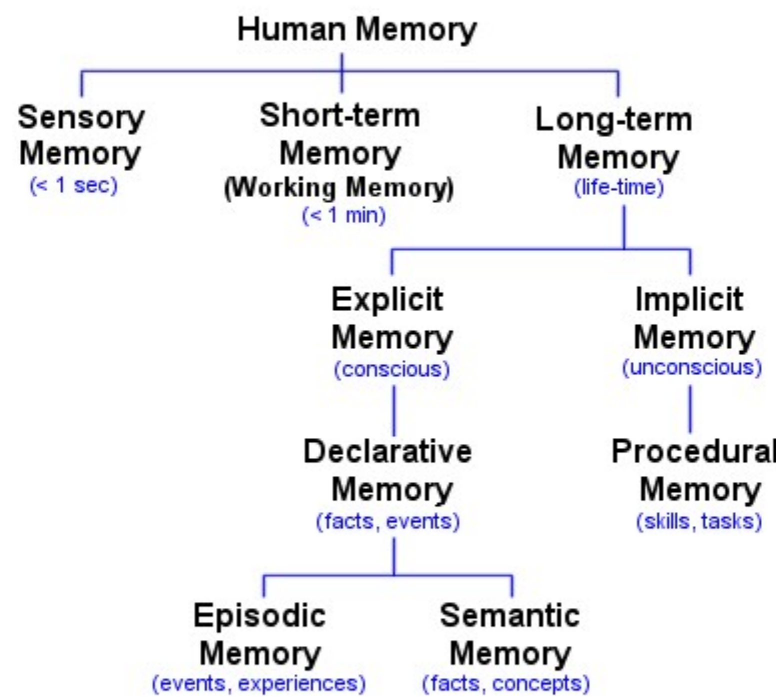

# Seven Sins of Memory

| Sin | Description | Example |
|---|---|---|
| ***Transience*** | Decreasing accessibility of memory over time | Simple forgetting of long-past events |
| ***Absent-Mindedness*** | Lapses of attention that result in forgetting | Forgetting location of car keys |
| ***Blocking*** | Information is present but temporarily accessible | Tip-of-the-tongue |
| ***Misattribution*** | Memories are attributed to an incorrect source | Confusing a dream for a memory |
| ***Suggestibility*** | Implanted memories about things that never occurred | Leading questions produce false memories |
| ***Bias*** | Current knowledge and beliefs distort our memories of the past | Recalling past attitudes in line with current attitudes |
| ***Persistence*** | Unwanted recollections that we can never forget | Traumatic war memories |

***Credit to Professor Daniel Schacter, Chair of Psychology at Harvard University***

# Case Study on Memory: Henry Molaison (HM)

# Henry Molaison and Brenda Milner

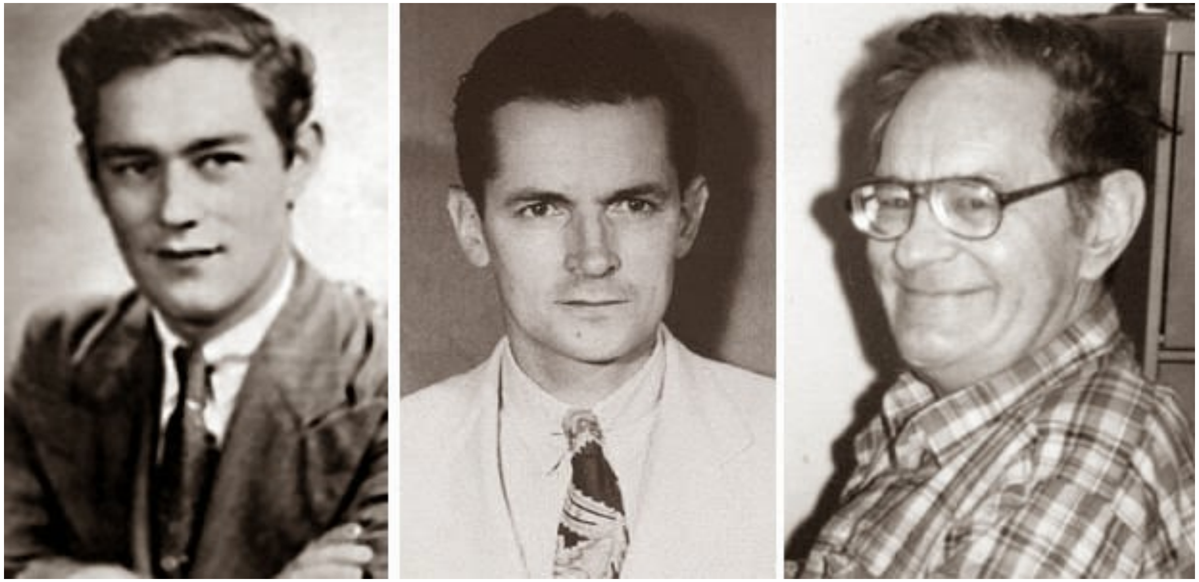

Understanding the Human Brain: A Lifetime of Dedicated Pursuit

***Thank You, Professor Brenda Milner***

# Case Study on Memory: Henry Molaison

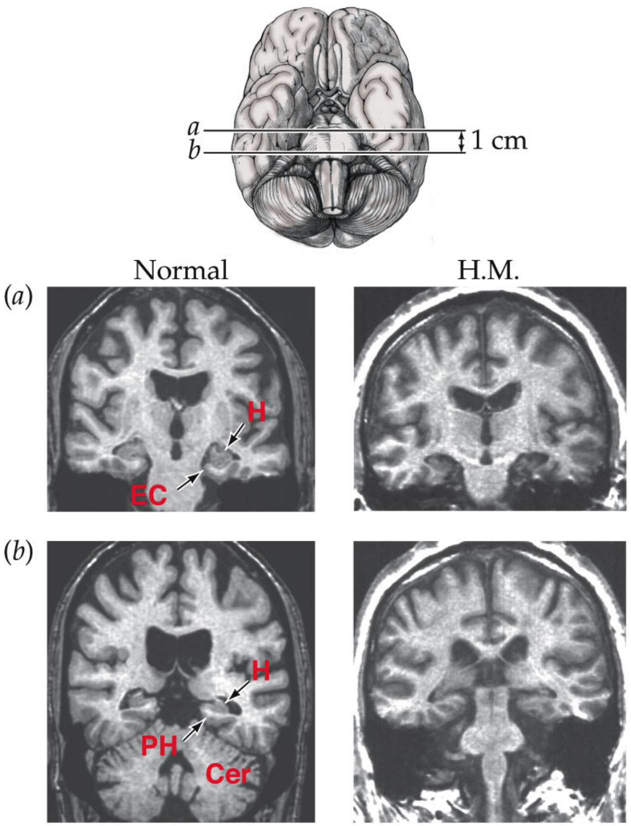


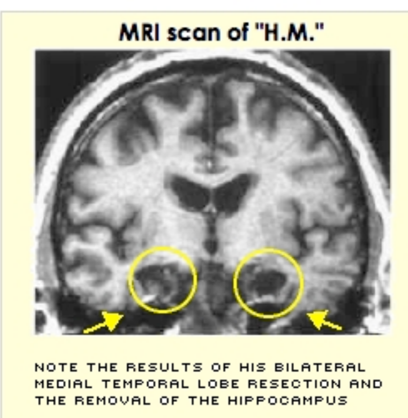


**Anterograde Amnesia:** Problems of learning new facts

- Specific to episodic memories
- Procedural memories intact
- Implicit memory performance intact
- Verbal learning disrupted

Damage to the hippocampus or to regions that supply its inputs and receive its outputs causes anterograde amnesia *as evidenced by Henry Molaison Case Study*

# Role of the Hippocampus

1. **Formation of new episodic memories**

   *Anterograde amnesia (HM)*

2. **Cognitive Map**

   *Place Cells in Rats*

   *Spatial Attention Cells in Monkeys*

3. **Configurable Association Theory**

   *Rats with hippocampus lesions are impaired on tasks requiring them to recognize cue configurations*

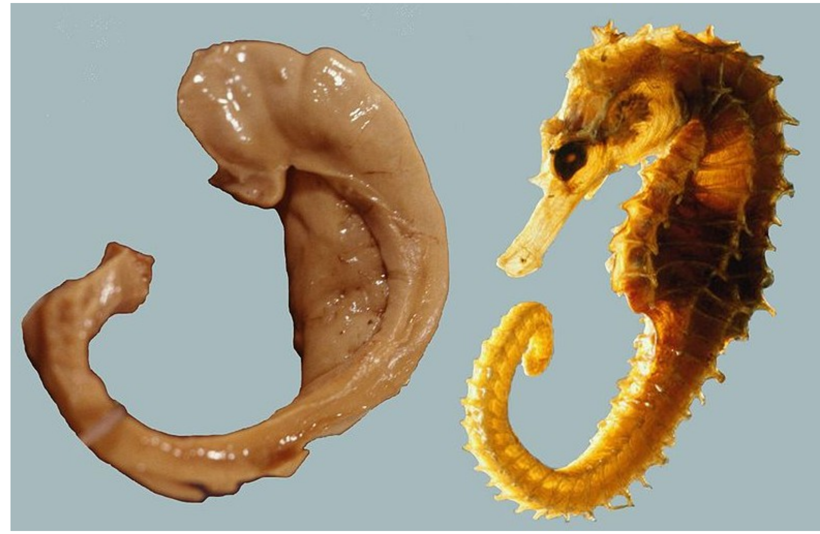

*Hippocampus means "seahorse"*

# Place Cells in Rats

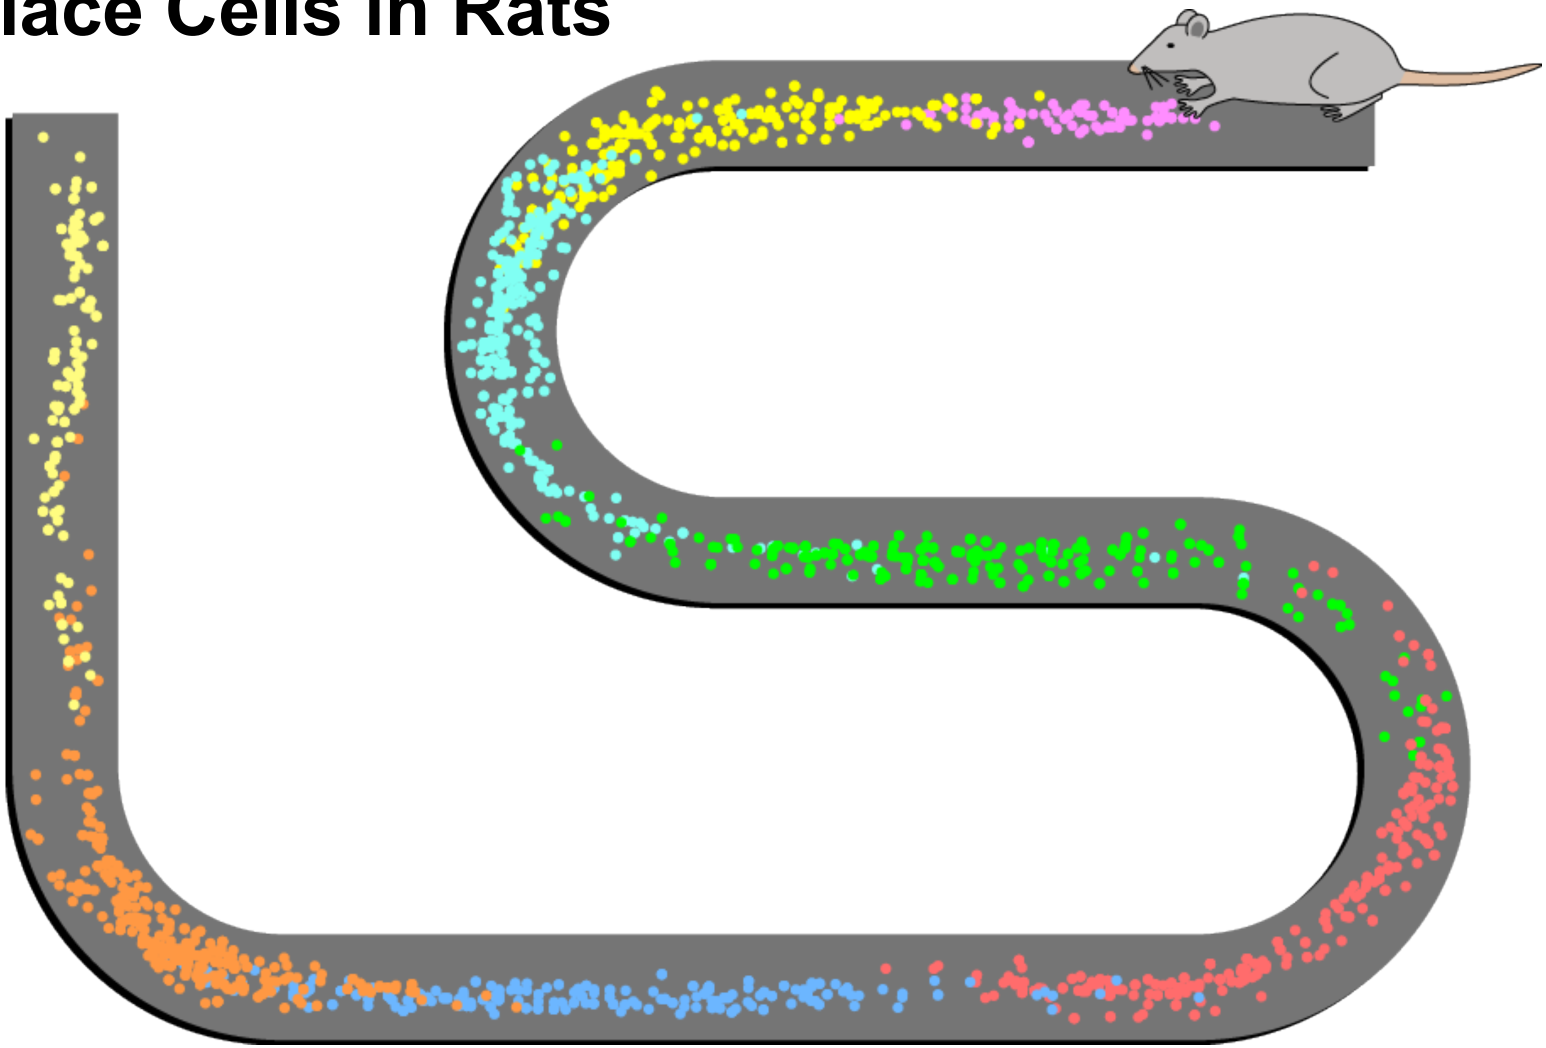

**Place Cell:** A type of pyramidal neuron within the hippocampus that becomes active when an animal enters a particular place in its environment (***place field***)

# Case Study on Memory: Alzheimer's Disease

# Alzheimer's Disease (AD)

- Cortical and Progressive Dementia

  **Pre-Clinical AD**

  - Signs of AD are first seen in entorhinal cortex and then proceed to the hippocampus
  - Affected regions begin to shrink as nerve cells die
  - Changes can begin 10 to 20 years before symptoms appear
  - **Memory Loss** is the **first** sign of AD

  **Mild to Moderate AD**

  - AD spreads through the brain and the cerebral cortex begins to shrink as more and more neurons die

  **Severe AD**

  - Extreme shrinkage, patients are completely dependent on others for care.
  - Death from pneumonia and other infections

# Alzheimer's Disease (AD)

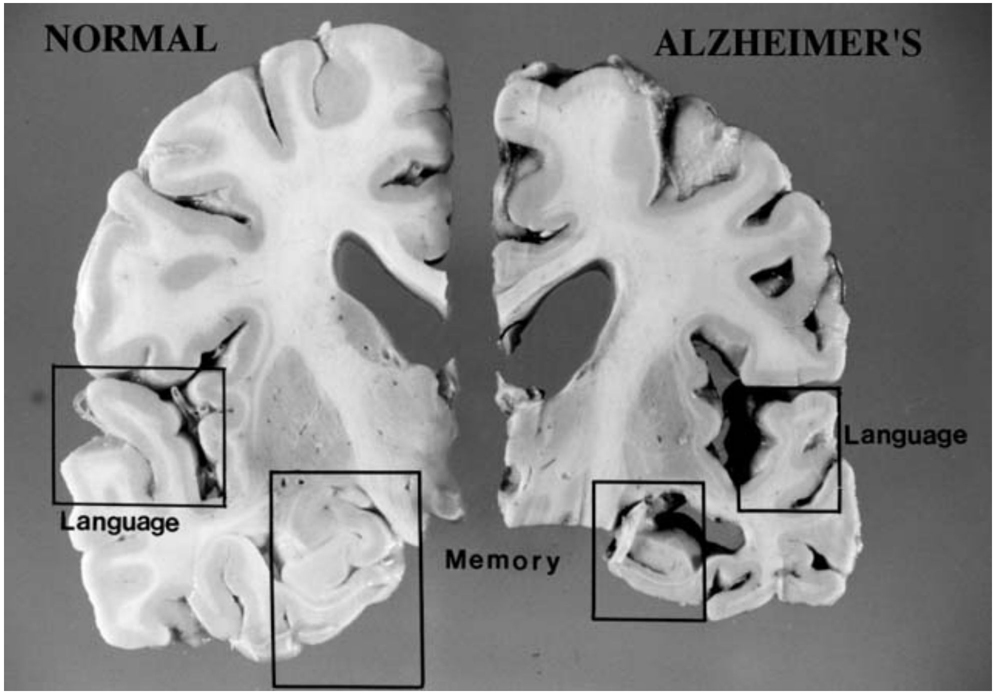


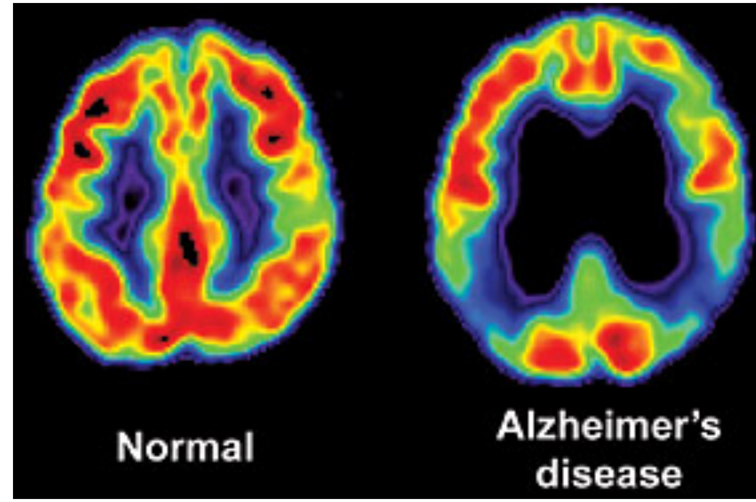


1. Extreme Shrinkage of Cerebral Cortex
2. Extreme Shrinkage of Hippocampus
3. Severely enlarged ventricles

**Positron Emission Tomography (PET):**
Uses a small amounts of radiotracers (analog of glucose, *fludeoxyglucose*) to evaluate organ and tissue function

> *Intense labeling of organs that use glucose extensively*

# Alzheimer's Disease (AD)

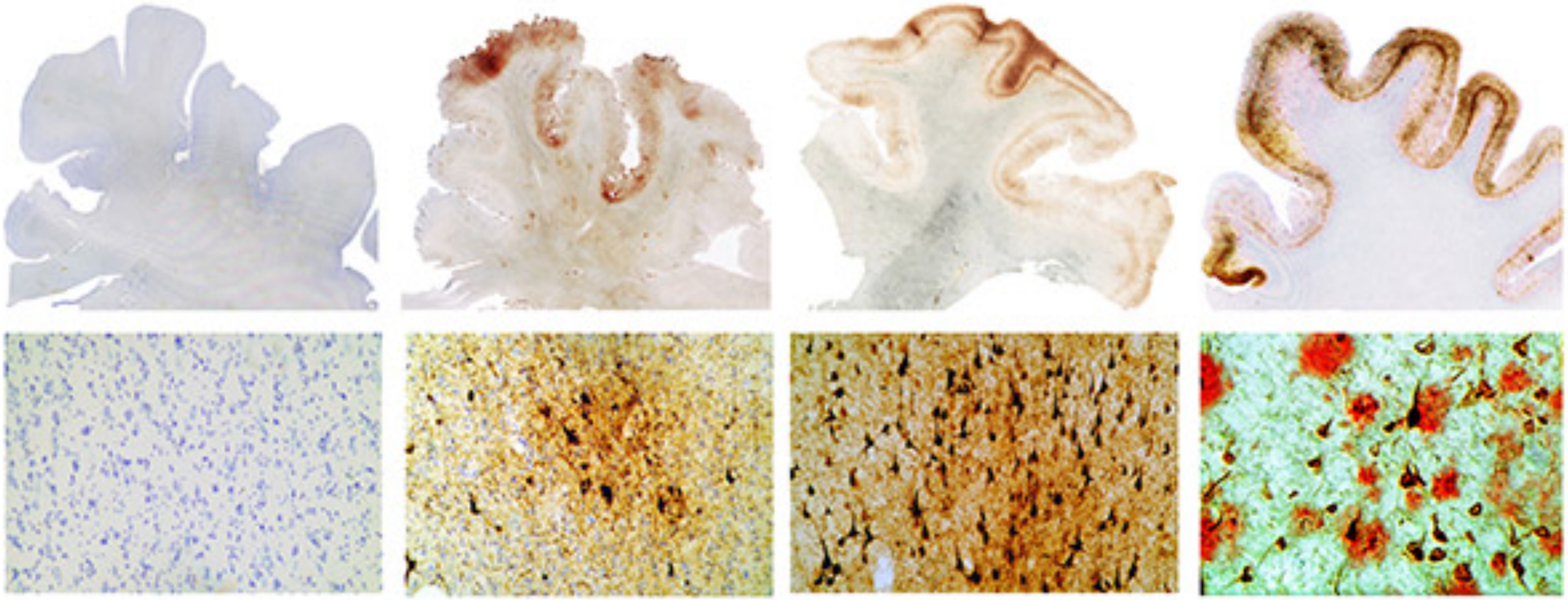

**Neurofibrillary Tangles (NFT):** aggregates of hyper-phosphorylated tau protein (changes in cytoskeleton)

**B-Amyloid Plaques:** Peptides of 36-43 amino acids that aggregate as plaques (incorrectly folded proteins)

# Case Study on Memory: Jill Price

# Case Study on Memory: Jill Price

# Jill Price

Hyperthymesia

- Condition of possessing an extremely detailed autobiographical memory
- Studied by three UCI neurobiologists: Elizabeth Parker, Larry Cahill and James McGaugh
- Attributed it to two defining characteristics:
  - Spending an excessive amount of time thinking about one’s past
  - Displaying an extraordinary ability to recall specific events from one’s past
- Causes:

  Psychological: information coded is semantic, so semantic clues are used in retrieval
    - Once memory is retrieved, it is episodic and follows a spreading activation model

  Biological: temporal lobe and caudate nucleus were both enlarged – can be attributed to atypical neural development

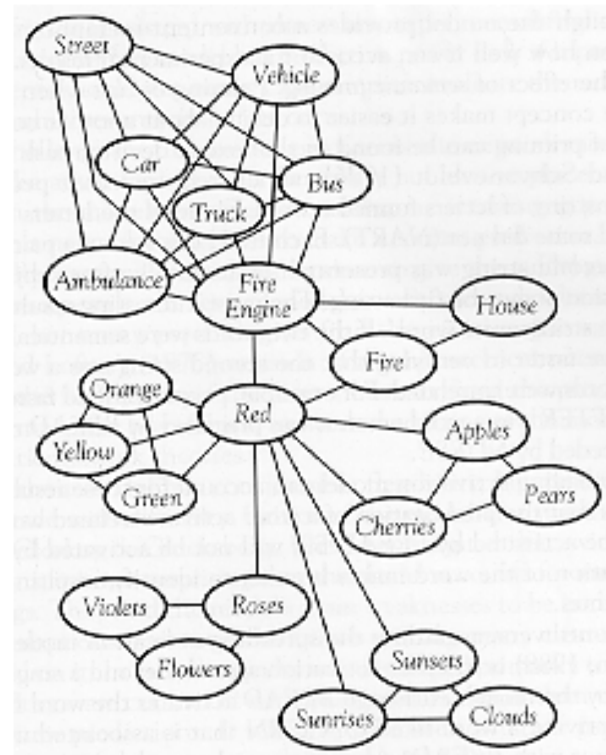

# Brain Regions for Memory

Case studies have taught us what brain regions are involved in encoding memories:

| **Memory** | **Brain Regions Involved** |
|---|---|
| Recalling pictures | Right prefrontal cortex and parahippocampal cortex of both hemispheres |
| Recalling words | Left prefrontal cortex and left parahippocampal cortex are activated |
| Consolidation of Memory | Hippocampus |
| Storage of Long-Term Memory | Cerebral Cortex (near where memory was first processed and held in short-term memory) |

Summary:
**1. Pre-Frontal Cortex**
**2. Hippocampus**
**3. Cerebral Cortex**
**4. Amygdala** (Memory Modulation and part of **Temporal Lobe**)

# Memories are made of this

*"He was still too young to know that the heart's memory eliminates the bad and magnifies the good, and that thanks to this artifice we manage to endure the burden of the past."*
***Gabriel Garcia Marquez***

**Pyramidal Neurons:** found in the cerebral cortex, the hippocampus, and the amygdala
***Pyramidal neurons are the primary excitation units of the mammalian pre-frontal cortex and the cortico-spinal tract.***

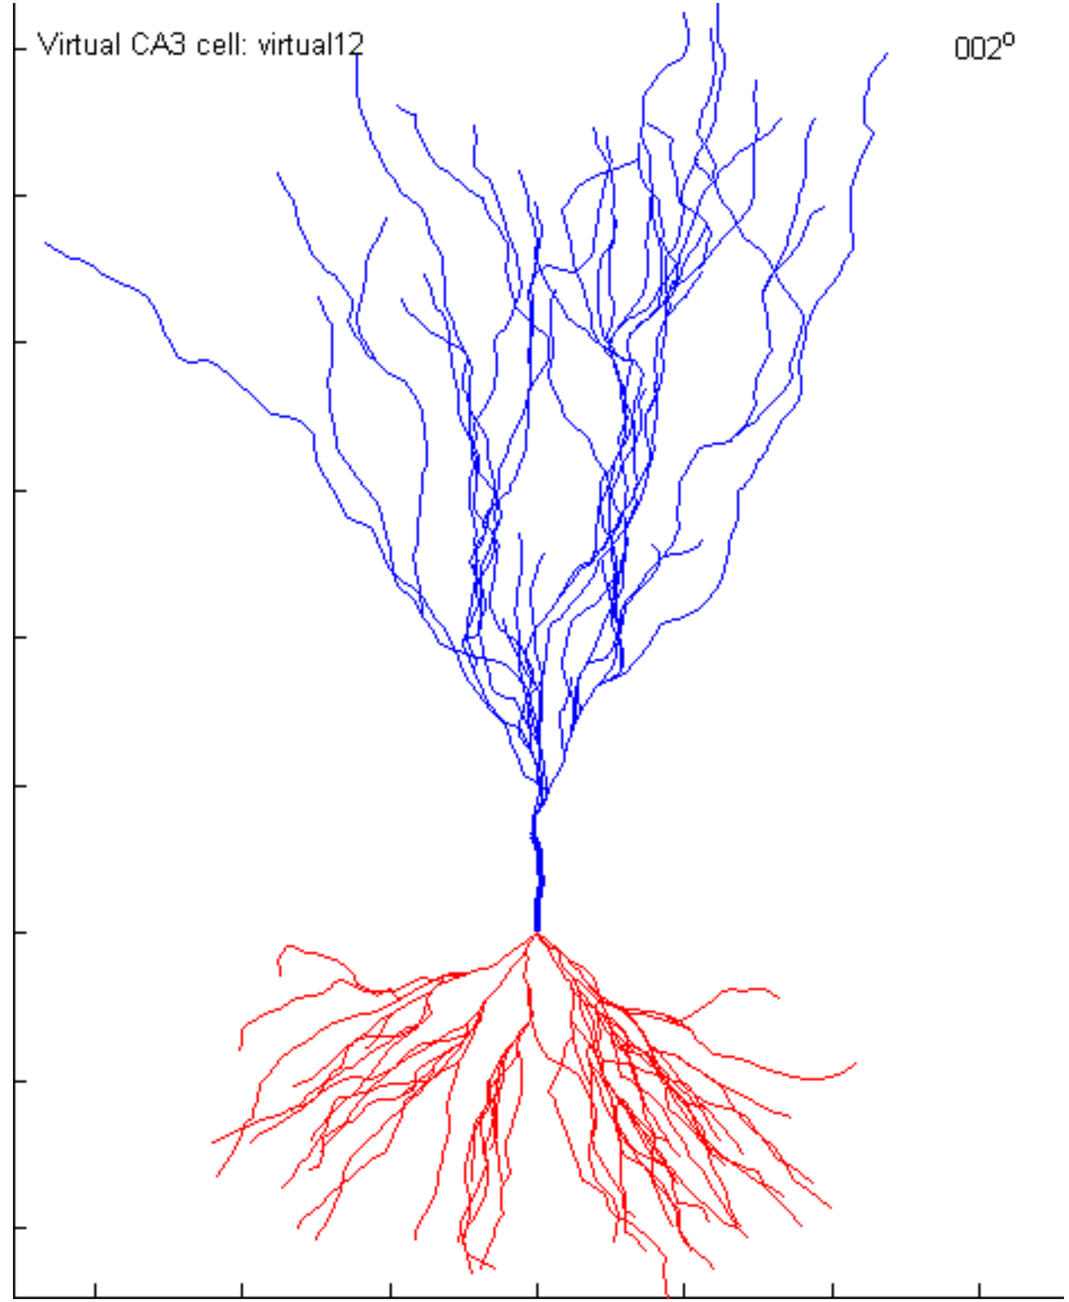

# What creates a memory in the brain?

Many biophysical variables are involved, including:

- *Neural voltage*
- *Synaptic activation, strength, and connectivity*
    - *How does a chemical synapse work?*
- *Pre-synaptic vesicles*
    - *What variable in regards to action potentials controls the amount of neurotransmitter release on the pre-synaptic neuron?*
- *Phosphorylation levels*
    - *What does a phosphate group look like?*
- *mRNA concentrations*
- *Transcriptional regulation*
- *Neuro-modulatary signals*
    - *What part of the brain is involved in modulation of memory?*
- *Glia*

# Man v. Machine

How can we compare the forms of memory in humans and computers?

Structures of Memory:

- **Random Access Memory (RAM) and the Hippocampus**
- **The Hard Drive and the Cerebral Cortex**
- **Central Processing Unit (CPU) Cache and Neuronal Network Attractors**

Features of Memory:

- **Protected Memory and Explicit Memory**
- **Memory Swapping and Writing**

# Random Access Memory (RAM) and the Hippocampus


**Random Access Memory:** place in a computing device where the OS, application programs and data current in use are kept so they can be reached by the device's processor

***Data remains in the RAM as long as the computer is running. If you turn off a computer, the RAM is gone.***

RAM provides a quick access to a memory in any location.

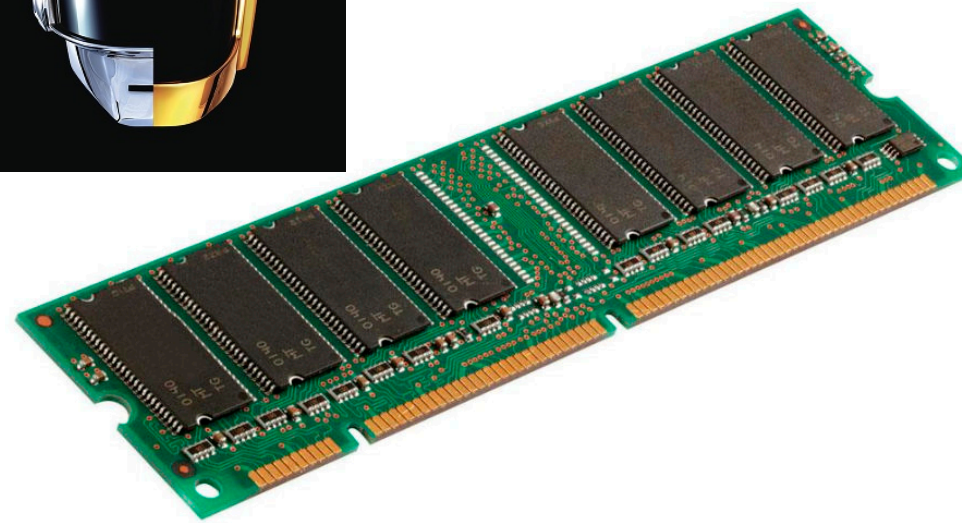

Humans have a form of RAM: **Short-Term Memory** (used to recall immediate actions). Both share **structural homogeneity.**

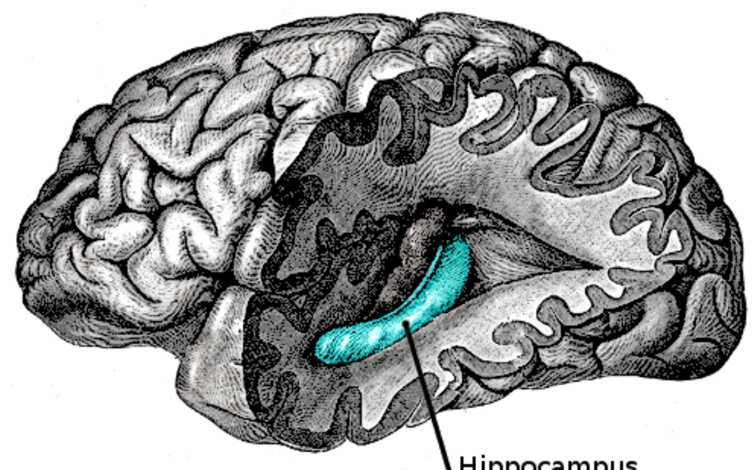

# The Hard Drive and the Cerebral Cortex

In addition to RAM, a complimentary form of storage for long-term memory in computers is **the Hard Drive (HD).** A hard drive has lower bandwidth (100 Mbytes/s) but can store much more (500 Gbytes). **This is where all of your data and programs are located.**

**Bandwidth:**
[1](Computer) Range of frequencies used for transmitting a signal
[2](Human) Energy or mental capacity required to deal with a situation

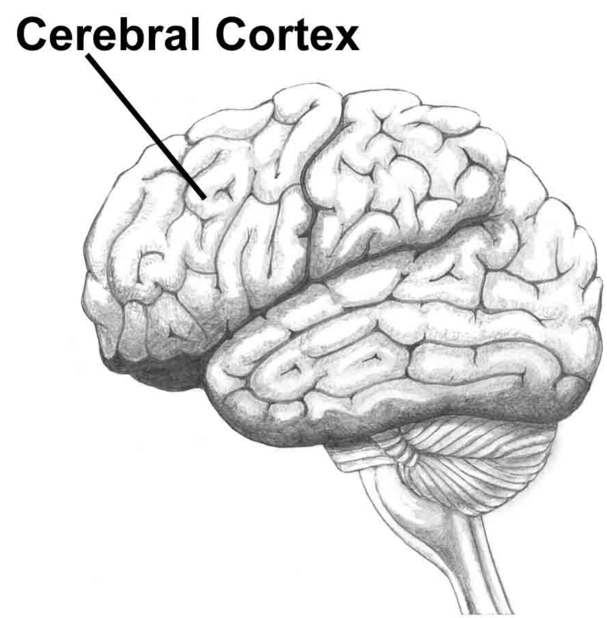


**Cerebral cortex** is the presumed site of human long-term memory. Recollecting old memories varies in timescale for both the cortex and HD (read/write speed depends on where memory is stored).

# Central Processing Unit (CPU) Cache and Neuronal Network Attractors


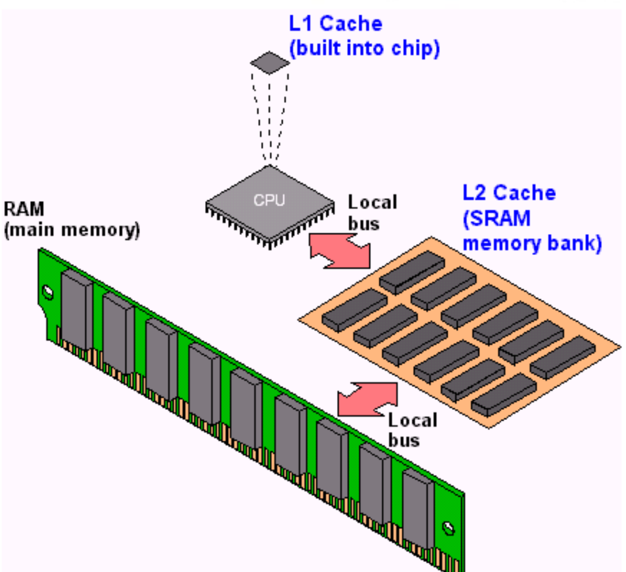


**CPU Cache:** a smaller, faster memory which stores copies of data from frequently used main memory locations

**How Caching Works (Transfer of RAM)**
CPU caches are small pools of memory that store information the CPU is most likely to need next.
**L1 Cache** is built into microprocessor chip itself and performs **cache hits.**
**A cache miss,** however, means the CPU must find data elsewhere. The **L2 Cache** is used for this (It is slower but larger than L1 Cache).

***A CPU Cache is the main difference between a computer and a human. Our brains do not have multiple levels of storage working at distinct speeds. Memory storage is distributed all over cortex (a giant cache).***

# Modeling Memory: Attractor Neural Networks

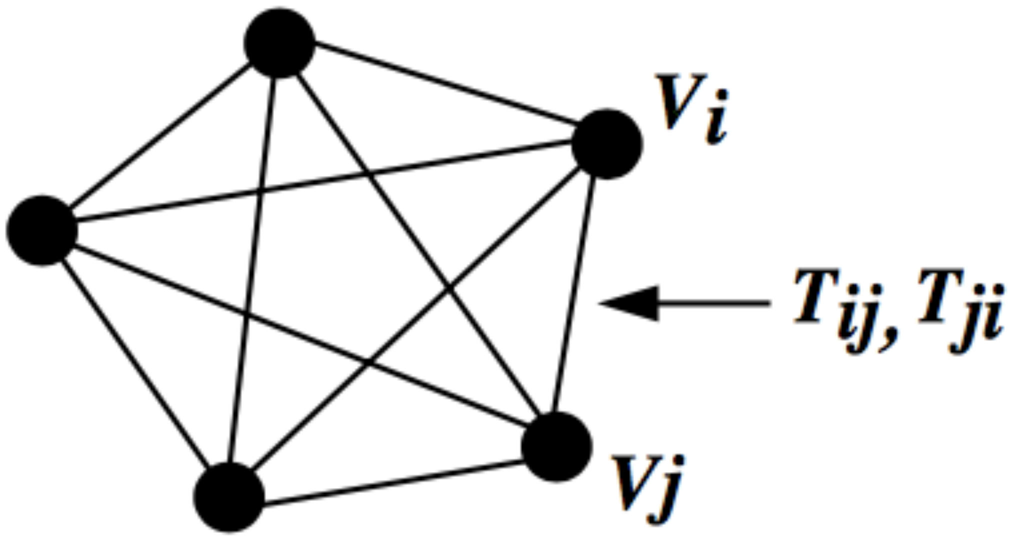


***$V_i$*: *denotes activity state of* neuron i *in the network***

***$T_{ij}$*: *denotes strength of the connection from* neuron j *to* neuron i**

**Attractor Neural Network:** recurrent neural network with symmetric connections that act in two ways

$$V_i \rightarrow 1 \quad \textit{if} \quad \sum_{j \neq i} T_{ij} V_j > 0$$

$$V_i \rightarrow -1 \quad \textit{if} \quad \sum_{j \neq i} T_{ij} V_j < 0$$

Each neuron has two states, **+1 and -1**, and changes its state, $V_i$, according to rule on left

# Features of Memory: Protected Memory and Explicit Memory

**Explicit Memory:** Memory used consciously
**Implicit Memory:** Memory not necessarily used consciously
*(Example: We know how to walk but we don’t have to think to walk)*
***This separation of between explicit and implicit memory is done at neuronal level and happens in computers.***

**Memory Protection:** way to control memory access rights on a computer that is built-into an OS
*When you write a program that needs some memory for its behavior, the OS allocates some space into it. Why? To prevent the program from accessing any other memory and rewriting it*

***Distinguishing our memories is a way to protect them.***

# Features of Memory: Memory Swapping and Writing

***What is memory swapping?***

When a program asks for memory then what is available in the RAM, the OS uses the hard drive (HD) to allocate space needed. This is known as **memory swapping.**

Since hard drives are slower than RAM, this leads to a great loss in performance.

***This is why MATLAB programs can slow down. Your program cannot tell if its data is being stored on the hard drive or the RAM. In other words, memory swapping is not transparent.***

***Do humans do memory swapping?***

No, the human brain was well-designed to work within its memory limits. However, we compensate for our memory by taking notes *(we swap data to a piece of paper using a pen).*

# Memory Storage System

Based on what we see in computers and humans, we can define features we desire in a **memory storage system:**

1. States that can persist over time
2. Adequate storage capacity (can hold a great number of states)
3. Different inputs to be remembered should trigger ***persistence*** of different memory states
4. Memory states are robust to noise
   *What are examples of noise in the brain?*
   Signals can fail and neurons can spike stochastically.
5. Memories are retrievable given appropriate cues

*Which feature would differ between short-term memory and long-term memory?*

# Creating a False Memory in the Hippocampus

*by Steve Ramirez, Xu Liu, Pei-Ann Lin, Junghyup Suh, Michele Pignatelli, Roger L. Redondo, Tomás J. Ryan, and Susumu Tonegawa*

*Science*
*Volume 341(6144):387-391*
*July 26, 2013*

**Fig. 1 Activity-dependent labeling and light-activation of hippocampal neurons, and the basic experimental scheme.(A) The c-fos-tTA mice were bilaterally injected with AAV9-TRE-ChR2-mCherry and implanted with optical fibers targeting DG. (B) While on Dox, exploration of a novel context did not induce expression of ChR2-mCherry.**

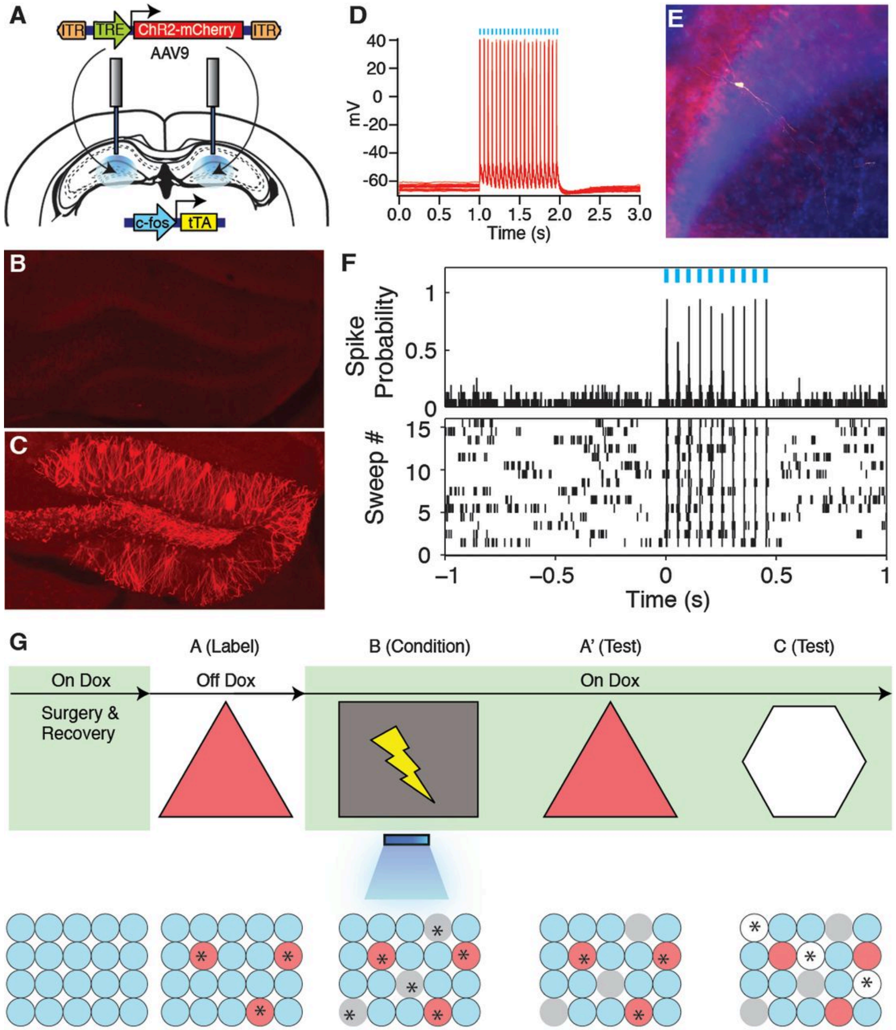


Steve Ramirez et al. Science 2013;341:387-391

Published by AAAS

**Fig. 2 Creation of a false contextual fear memory.(A to E) c-fos-tTA mice injected with AAV9-TRE-ChR2-mCherry in the DG were taken off Dox and exposed to context A in order to label the activated cells with mCherry (red), then put back on Dox and exposed to the same context A [(A) and (C)] or a novel context C [(B) and (D)] 24 hours later so as to let activated cells express c-Fos (green).**

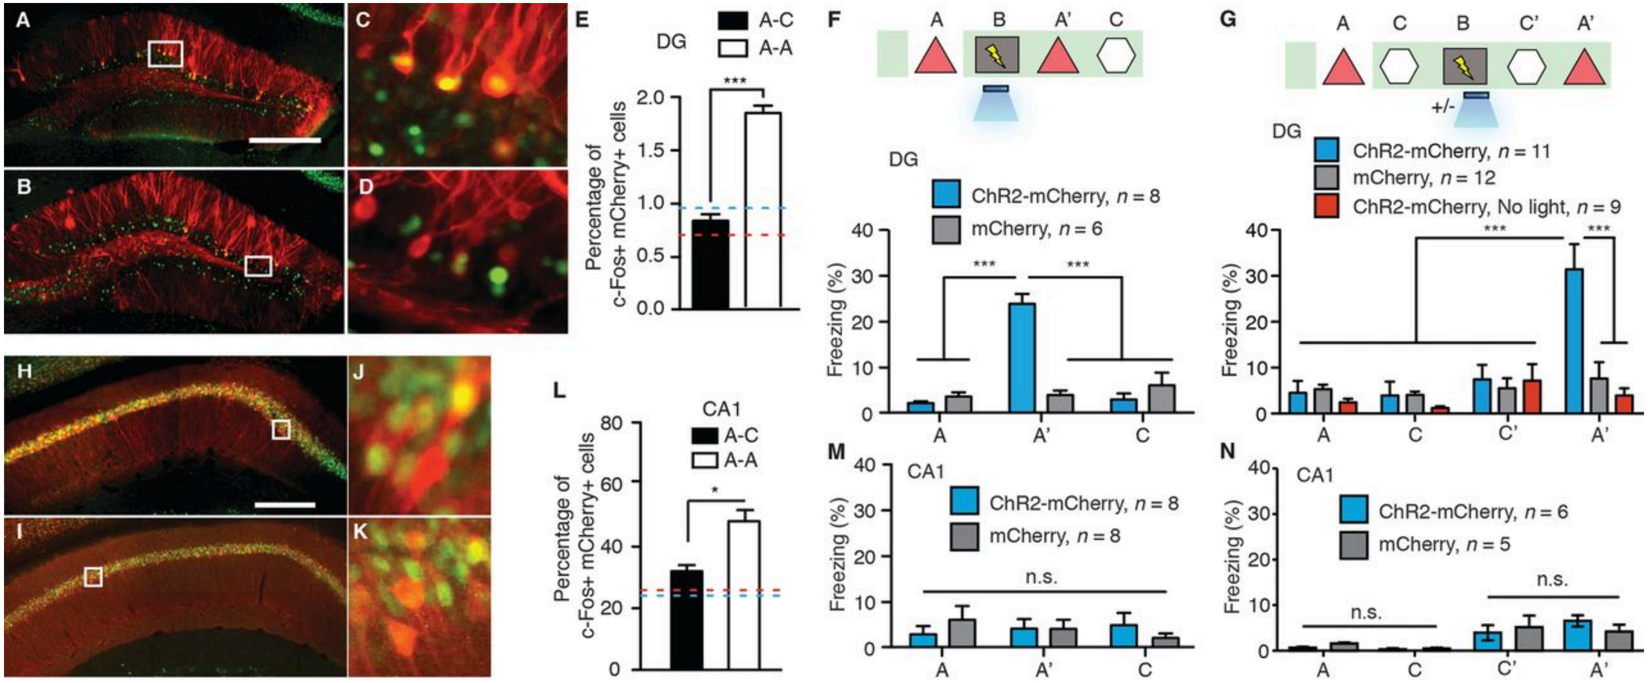


Steve Ramirez et al. Science 2013;341:387-391

Published by AAAS

**Fig. 3 The false and genuine fear memories interact with each other, and both recruit the amygdala.(A) Animals that underwent the behavioral protocol shown in Fig. 2G were reexposed to context B, and the freezing levels were examined both in the absence and presence of light stimulation (n = 11 subjects for ChR2-mCherry group and n = 9 subjects for ChR2-mCherry, no-light group; *P = 0.027; ***P < 0.001; #P = 0.034, two-way ANOVA with repeated measures followed by Bonferroni post-hoc test).**

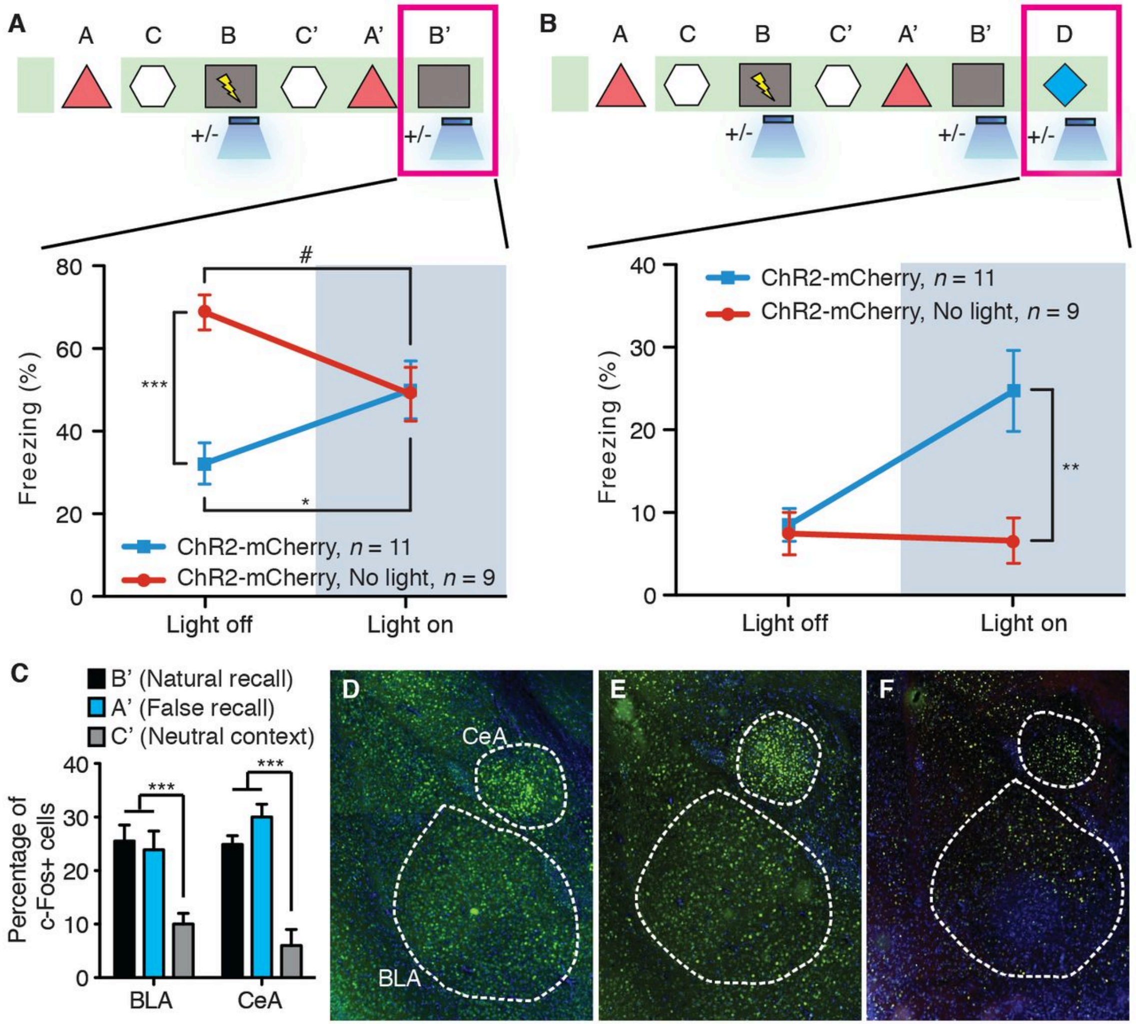


Published by AAAS

## Fig. 4 The false memory supports active fear behavior.(A) The scheme for conditioned place-avoidance paradigm.

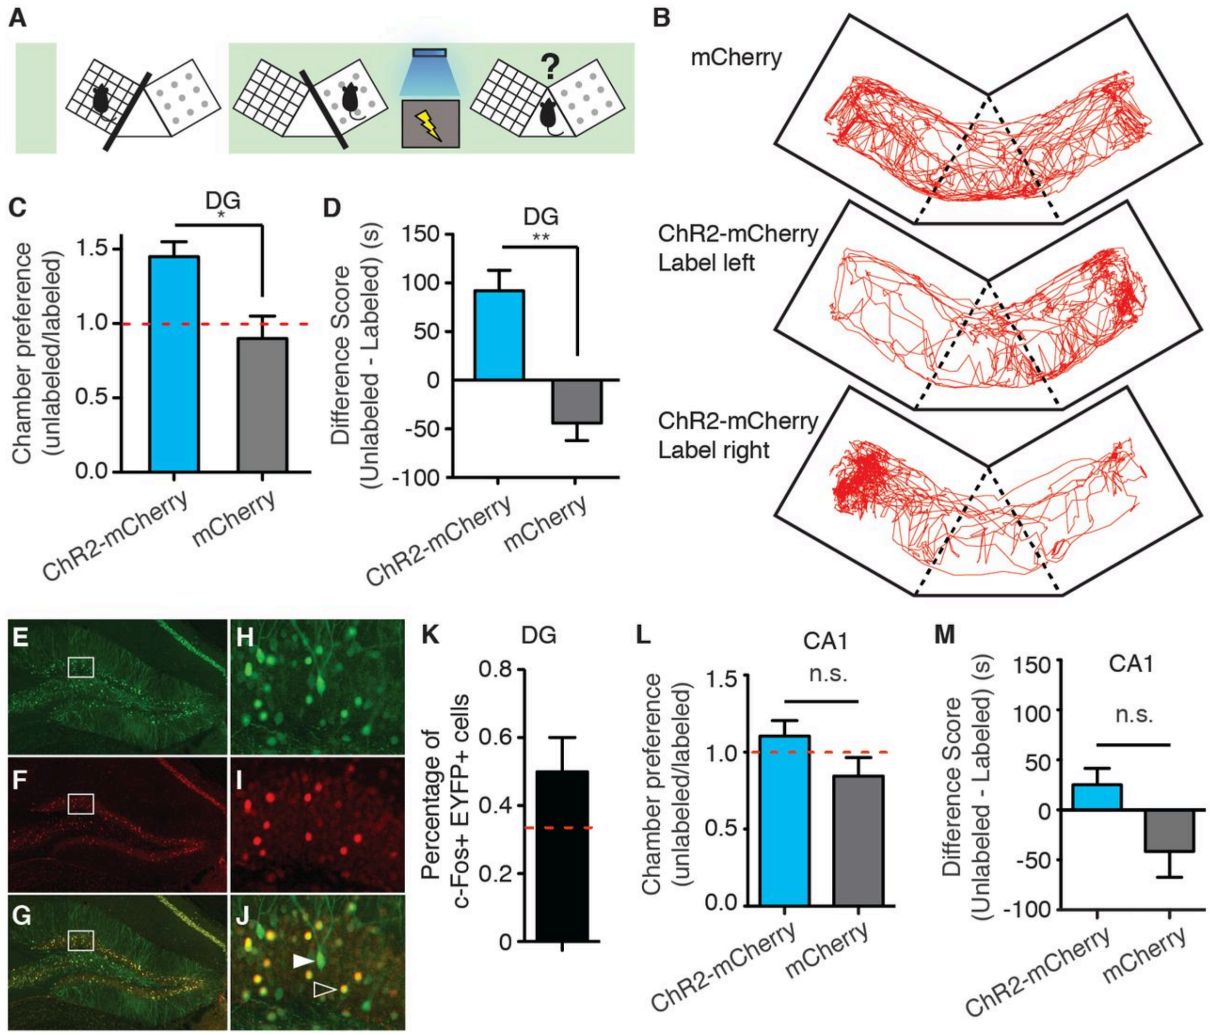


Steve Ramirez et al. Science 2013;341:387-391

# Next Time: Glia